arXiv:1611.02460v1 [cs.DM] 8 Nov 2016

On coalescence time in graphs

When is coalescing as fast as meeting?

Varun Kanade
University of Oxford
varunk@cs.ox.ac.uk

Frederik Mallmann-Trenn
École normale supérieure
Simon Fraser University
mallmann@di.ens.fr

Thomas Sauerwald*
University of Cambridge
thomas.sauerwald@cl.cam.ac.uk

Abstract

Coalescing random walks is a fundamental stochastic process, where a set of particles perform independent discrete-time random walks on an undirected graph. Whenever two or more particles meet at a given node, they merge and continue as a single random walk. The *coalescence time* is defined as the expected time until only one particle remains, starting from one particle at every node. Despite recent progress such as Cooper et al. [14] and Berenbrink et al. [7], the coalescence time for graphs such as binary trees, d-dimensional tori, hypercubes and more generally, vertex-transitive graphs, remains unresolved.

We provide a powerful toolkit that results in tight bounds for various topologies including the aforementioned ones. The meeting time is defined as the worst-case expected time required for two random walks to coalesce. As a general result, we establish that for graphs whose meeting time is only marginally larger than the mixing time (a factor of $\log^2 n$), then the coalescence time of n random walks equals the meeting time up to constant factors. This upper bound is complemented by the construction of a graph family demonstrating that this result is the best possible up to constant factors. For almost-regular graphs, we bound the coalescence time by the hitting time, resolving the discrete-time variant of a conjecture by Aldous for this class of graphs. Finally, we prove that for any graph the coalescence time is bounded by $O(n^3)$. By duality, our results give bounds on the voter model.

keywords— coalescing time, meeting time, random walks, voter model

*Supported by the ERC Starting Grant (DYNAMIC MARCH).

Contents

1 Introduction **1**
- 1.1 Contributions . . . 2
- 1.2 Proof Ideas and Technical Contributions . . . 5
- 1.3 Discussion and Future Work . . . 9

2 Notation and Preliminaries **9**

3 Bounding t_{coal} for large $t_{\text{meet}}/t_{\text{mix}}$ **10**
- 3.1 Stochastic Process . . . 10
- 3.2 Meeting Time Distribution Prior to t_{meet} . . . 13
- 3.3 Upper Bound - Proof of Theorem 1.1 . . . 14
- 3.4 Lower Bound - Proof of Theorem 1.2 . . . 18

4 Bounding $t_{\text{coal}} = O(t_{\text{hit}})$ for Almost-Regular Graphs **26**
- 4.1 Concentration Inequalities for Random Walks . . . 27
- 4.2 Reducing the Walks from n to $O(\log^3 n)$ in $O(t_{\text{hit}})$ . . . 29
- 4.3 Reducing the Walks from $\log^4 n$ to $(\Delta/d)^{O(1)}$ in $O(t_{\text{hit}})$ . . . 31
- 4.4 Bounding t_{coal} in terms of t_{hit} . . . 39
- 4.5 Proof of Theorem 1.3 . . . 39
- 4.6 Conjecture and a Possible Improvement for Non-Regular Graphs . . . 39

5 Bounding $t_{\text{coal}} \in [\Omega(\log n), O(n^3)]$ **40**
- 5.1 General Upper Bound $t_{\text{coal}} = O(n^3)$ . . . 40
- 5.2 General Lower Bound $t_{\text{coal}} = \Omega(\log n)$ . . . 41
- 5.3 Proof of Theorem 1.5 . . . 41

6 Bounding t_{meet} and Applications **42**
- 6.1 Relating Meeting Time to t_{mix} and $\frac{1}{1-\lambda_2}$ . . . 42
- 6.2 Relating Meeting Time to t_{hit} . . . 48
- 6.3 Proof of Theorem 1.4 . . . 50

7 Applications to Special Graphs **50**
- 7.1 2-Dimensional Grids/Tori and Paths/Cycles . . . 50
- 7.2 d-Dimensional Grids and Tori, $d \geqslant 3$ . . . 50
- 7.3 Hypercubes . . . 51
- 7.4 (Regular) Expanders . . . 51
- 7.5 Real World Graph Models . . . 51
- 7.6 Binary Trees . . . 52
  - 7.6.1 Bounds on the t-step probabilities . . . 52
  - 7.6.2 Establishing the Lower Bound on the Meeting Time . . . 54
- 7.7 Star . . . 56

A Basic Results about Markov Chains **58**

1 Introduction

Coalescing random walks is a fundamental stochastic process on *connected* and *undirected* graphs. The process begins with particles on some subset of the nodes in the graph. At discrete time-steps, every particle performs one step of a random walk.[1] Whenever two or more particles arrive at the same node at the same time-step, they merge into a single particle and continue as a single random walk. The *coalescence time* is defined as the first time-step when only one particle remains. The coalescence time depends on the number and starting positions of the particles.

When starting with two particles, the coalescence time is referred to as the *meeting time*. Let t_{meet} denote worst-case expected meeting time over all pairs of starting nodes and let t_{coal} denote the expected coalescence time starting from one particle on every node. It is clear that $t_{\text{meet}} \leqslant t_{\text{coal}}$; as for an upper bound, it can be shown that $t_{\text{coal}} = O(t_{\text{meet}} \log n)$, where n is the number of nodes in the graph. The main idea used to obtain the bound is that the number of surviving random walks halves roughly every t_{meet} steps. A proof of the result appears implicitly in the work of Hassin and Peleg [22].

Aldous [1] showed in *continuous-time* that the meeting time is bounded by the maximum hitting time, *i. e.,* $t_{\text{hit}} := \max_{u,v} t_{\text{hit}}(u, v)$, where $t_{\text{hit}}(u, v)$ denotes the expected time required to hit v starting from vertex u.We observe that the result of Aldous also holds in discrete time. Thus, this gives a bound of $O(t_{\text{hit}} \log n)$ for the coalescing time; however, in general $O(t_{\text{hit}})$ may be a loose upper bound on t_{meet}. In recent work, Cooper *et al.* [14] provide results that are better than $O(t_{\text{meet}} \log n)$ for several interesting graph classes, notably expanders and power-law graphs. They show that $t_{\text{coal}} = O((\log^4 n + \|\pi\|_2^{-2}) \cdot (1 - \lambda_2)^{-1})$, where λ_2 is the second largest eigenvalue of the transition matrix of the random walk and π is the stationary distribution. Berenbrink *et al.* [7] show that $t_{\text{coal}} = O(m/(d_{\min} \cdot \Phi))$, where m is the number of edges, $d_{\min}$ is the minimum degree and Φ is the conductance. Their result improves on that of Cooper *et al.* for certain graph classes, *e. g.,* cycles.

Despite the recent progress due to Cooper *et al.* [14] and Berenbrink *et al.* [7], for many fundamental graphs such as the binary tree, hypercube and the (d-dimensional) torus, the coalescing time in the discrete setting remains unsettled. We provide a rich toolkit allowing us to derive tight bounds for many graphs including all of the aforementioned ones. One of our main results establishes a relationship between the ratios $t_{\text{coal}}/t_{\text{meet}}$ and $t_{\text{meet}}/t_{\text{mix}}$, where t_{mix} denotes the mixing time.[2] In particular, the result shows that if $t_{\text{meet}}/t_{\text{mix}} = \Omega(\log^2 n)$, then $t_{\text{coal}} = O(t_{\text{meet}})$; however, we also provide a more fine-grained tradeoff. For almost-regular graphs, we bound the coalescence time by the hitting time. For vertex-transitive graphs we show that the coalescence time, the meeting time, and the hitting time are equal up to constants. Finally, we prove that for any graph the coalescence time is bounded by $O(n^3)$. This last result was mistakenly stated, as a side remark, to be a simple consequence from the main result of Cooper *et al.* [14]. Prior to that, [22] had shown a worst-case upper bound of $O(n^3 \log n)$. We also give worst-case upper and lower bounds on the meeting time and coalescence time that are tight for general graphs and regular (or nearly-regular) graphs.

In the process of establishing bounds on the coalescence time, we developed techniques to give tight bounds on the meeting time. We applied these to various topologies such as the binary tree, torus and hypercube. We believe that these techniques might be of more general interest.

The process of coalescing random walks, henceforth also referred to as the coalescence process,

[1] Throughout this paper, we use random walk and particle interchangeably, assuming that every random walk has an identifier.

[2] The *mixing time* is the first time-step at which the distribution of a random walk starting from an arbitrary node is close to the stationary distribution.

Graph	t_{mix}	t_{meet}		t_{coal}		t_{hit}
Binary tree	$\Theta(n)$	$\Theta(n \log n)$	*Thm. 1.3* & *Lem. 7.6*	$\Theta(n \log n)$	*Thm. 1.3* & *Lem. 7.6*	$\Theta(n \log n)$
Clique	$\Theta(1)$	$\Theta(n)$	[14, 7] & *Thm. 1.1*	$\Theta(n)$	[14, 7] & *Thm. 1.1*	$\Theta(n)$
Cycle	$\Theta(n^2)$	$\Theta(n^2)$	[7] & *Thm. 1.3*	$\Theta(n^2)$	[7] & *Thm. 1.3*	$\Theta(n^2)$
Rand. r-reg.	$\Theta(\log n)$	$\Theta(n)$	[13, 14, 7] & *Thm. 1.1*	$\Theta(n)$	[13, 14, 7] & *Thm. 1.1*	$\Theta(n)$
Hypercube	$\Theta(\log n \log \log n)$	$\Theta(n)$	*Thm. 1.3*	$\Theta(n)$	*Thm. 1.1*	$\Theta(n)$
Path	$\Theta(n^2)$	$\Theta(n^2)$	[7] & *Thm. 1.3*	$\Theta(n^2)$	[7] & *Thm. 1.3*	$\Theta(n^2)$
Star	$\Theta(1)$	$\Theta(1)$	folklore	$\Theta(\log n)$	[22] & *Prop. 3.3*	$\Theta(n)$
Torus ($d = 2$)	$\Theta(n)$	$\Theta(n \log n)$	*Thm. 1.3*	$\Theta(n \log n)$	*Thm. 1.3*	$\Theta(n \log n)$
Torus ($d > 2$)	$\Theta(n^{2/d})$	$\Theta(n)$	*Thm. 1.1*	$\Theta(n)$	*Thm. 1.1*	$\Theta(n)$

Table 1: A summary of bounds on the mixing, meeting, coalescence and hitting times for fundamental topologies for discrete-time random walks. All bounds on the mixing and hitting times appear directly or implicitly in [2].

is dual to the voter model, which is defined as follows. Initially, every node has a distinct opinion. At every round, each node chooses synchronously one of its neighbors at random and adopts that node's opinion. The *consensus time* is defined as the time it takes until only one opinion remains. The voting process viewed backwards is exactly the same as the coalescence process starting with a random walk on every node; thus, the coalescence time and consensus time have the same distribution. However, when considering the coalescence process starting with fewer than n random walks or the voter model with fewer than n distinct opinions, the duality no longer holds. Recently, there have been several studies on variants of the voter model *e. g.,* [15, 5, 6], however, the behavior of these processes is fundamentally different, which is why do not discuss them here.

The process of coalescing random walks has also been studied in *continuous time*; in this case, particles jump to a random neighboring node when activated according to a Poisson clock with mean 1. As Cooper and Rivera [12] recently pointed out *"It is however, not clear whether the continuous-time results apply to the discrete-time setting."* and to the best of our knowledge, there is no general way in which results in *continuous time* can be transferred to *discrete time* or vice versa, even when the random walks in discrete-time are *lazy*. In the continuous time setting, Cox [16] show that the coalescence time is bounded by $\Theta(t_{\text{hit}})$ for tori. Oliveira [29] showed that the coalescence time is $O(t_{\text{hit}})$ in general. In a different work, Oliveira [30] proved so-called mean field conditions, which are sufficient conditions for the coalescing process on a graph to behave similarly to that on a complete graph. His main result in [30], Theorem 1.2, implies that $t_{\text{coal}} = O(t_{\text{meet}})$ whenever $t_{\text{mix}} \cdot \pi_{\max} = O(1/\log^4 n)$. One of our main results, Theorem 1.1, implies $t_{\text{coal}} = O(t_{\text{meet}})$ whenever $t_{\text{mix}}/t_{\text{meet}} = O(1/\log^2 n)$. Notice that since $t_{\text{meet}} \geqslant 1/(\|\pi\|_2^2) \geqslant 1/\pi_{\max}$, our necessary condition seems more general—however, the results in Oliveira [30] also establish mean-field behavior, while ours are only concerned with the expected coalescence time, t_{coal}.

1.1 Contributions

In this work, we provide several results relating the coalescence and meeting times to each other and to other fundamental quantities related to random walks on undirected graphs. In particular, our focus is on understanding for which graphs the coalescence time is the same as meeting time, as we know that t_{coal} is always in the rather narrow interval of $[\Omega(t_{\text{meet}}), O(t_{\text{meet}} \cdot \log n)]$. As a consequence

of our results, we derive new and re-derive existing bounds on the meeting and coalescence times for several graph families of interest. These results are summarized in Table 1 and discussed in greater detail in Section 7. Formal definitions of all quantities used below appear in Section 2.

Our first main result relates t_{coal} to t_{meet} and t_{mix}. As already mentioned in the introduction, the crude bound $t_{\text{coal}} = O(t_{\text{meet}} \log n)$ is well-known. However, this bound is not in general tight, as demonstrated by our result below.

Theorem 1.1. *For any graph G, we have*

$$t_{\text{coal}} = O\left(t_{\text{meet}}\left(1 + \sqrt{\frac{t_{\text{mix}}}{t_{\text{meet}}}} \cdot \log n\right)\right)$$

Consequently, when $t_{\text{meet}} \geqslant t_{\text{mix}}(\log n)^2$, $t_{\text{coal}} = O(t_{\text{meet}})$.

The proof of Theorem 1.1 appears in Section 3. One interesting aspect about this bound is that it can be used to establish $t_{\text{coal}} = \Theta(t_{\text{meet}})$ even without having to know the quantities t_{meet} or t_{mix}. This flexibility turns out to particularly useful when dealing with random graph models for "real world" networks. Here, we establish (nearly-)tight and sublinear bounds in Section 7.5.

To see the power of this result, we remark that the main result of Cooper *et al.* [14, Theorem 1] can be reproven by combining [14, Theorem 2] with Theorem 1.1 (see Proposition 6.2).

Our next main result shows that the bound in Theorem 1.1 is tight up to a constant factor, which we establish by constructing an explicit family of graphs. Interestingly, for this family of graphs we also have $t_{\text{hit}} \gg t_{\text{meet}}$, thus showing that t_{hit} may be a rather loose upper bound for t_{coal} in some cases.

Theorem 1.2. *For any sequence $(\alpha_n)_{n\geqslant 0}$, $\alpha_n \in [1, \log^2 n]$ there exists a family of almost-regular graphs (G_n), with G_n having $\Theta(n)$ nodes and satisfying $\frac{t_{\text{meet}}}{t_{\text{mix}}} = \Theta(\alpha_n)$ such that*

$$t_{\text{coal}} = \Omega\left(t_{\text{meet}} \cdot \left(1 + \sqrt{\frac{t_{\text{mix}}}{t_{\text{meet}}}} \cdot \log n\right)\right).$$

The above two results show that that $t_{\text{meet}}/t_{\text{mix}}$ should be $\Omega(\log^2 n)$ to guarantee that $t_{\text{coal}} = O(t_{\text{meet}})$; a discussion on why $\log^2 n$ is the right factor can be found in Section 1.2.

A natural question is therefore whether in the case of structured sub-classes such as regular graphs, or vertex-transitive graphs, or special graphs such as grids, tori, binary trees, cycles, real-world (power-law) graphs, *etc.*, better bounds can be obtained through other methods. We provide results that are tight or nearly tight in several of these cases; some of these results were previously known using other methods, some are novel to the best of our knowledge.

Theorem 1.3. *The following hold for graphs of the stated kind*

(i) For any graph G,

$$t_{\text{coal}} = O\left(t_{\text{hit}} \cdot \log\log n\right).$$

(ii) For any graph G with maximum degree Δ and average degree d,

$$t_{\text{coal}} = O\left(t_{\text{hit}} + t_{\text{meet}} \cdot \log(\Delta/d)\right).$$

Hence for any almost-regular[3] graph G, $t_{\text{coal}} = O\left(t_{\text{hit}}\right)$.

[3] We call a graph *almost-regular* if $\deg(u) = \Theta(\deg(v))$ for all $u, v \in V$.

(iii) For any vertex-transitive G,

$$t_{\text{coal}} = \Theta\left(t_{\text{meet}}\right) = \Theta\left(t_{\text{hit}}\right).$$

(iv) In the case of binary trees, d-dimensional tori/grids, paths/cycles, expanders, hypercubes, random power law graphs[4] *we have* $t_{\text{coal}} = \Theta(t_{\text{meet}})$.

The proof of the first three statements of Theorem 1.1 appear in Section 4. We point out that since $t_{\text{meet}} = O(t_{\text{hit}})$ for any graph, Theorem 1.3 implies the bound $t_{\text{coal}} = O(t_{\text{hit}})$ not only for almost-regular graphs, but also for dense graphs graphs where $|E| = \Theta(n^2)$. This settles the discrete-time analogue of a conjecture by Aldous [2, Open Problem 14.13] for these graph classes.

Another natural question is to express t_{meet} or t_{coal} solely in terms of t_{mix}, the spectral gap $1 - \lambda_2$ or other connectivity properties of G. We derive several such bounds on t_{meet}, t_{hit} and t_{coal}.

As a by-product of our techniques, we also derive new bounds on t_{hit} and t_{cov}, the cover-time. The detailed results are given in Section 6, but we highlight the results for regular graphs here:

Theorem 1.4. *Let G be any almost-regular graph, and let*

$$\beta := \min\left\{t_{\text{mix}}, \frac{1}{1-\lambda_2} \cdot \log\left(\frac{1}{1-\lambda_2}\right)\right\},$$

where λ_2 is the second largest eigenvalue of the transition matrix P of a lazy random walk. Then it holds that $t_{\text{hit}} = O(n \cdot \sqrt{\beta})$. Consequently, $t_{\text{meet}} \leqslant t_{\text{coal}} = O(n \cdot \sqrt{\beta})$ and $t_{\text{cov}} = O(n \log n \cdot \sqrt{\beta})$.

We point out that so far the best possible bound on t_{coal} for regular graphs has been $t_{\text{coal}} = O(n/(1-\lambda_2))$ from [14]; actually, the same also follows by the bound based on the conductance in [7] and Cheeger's inequality. The best possible bound on t_{hit} (and t_{cov}) in terms of $1 - \lambda_2$, was $t_{\text{hit}} = O(n/(1-\lambda_2))$ and $t_{\text{cov}} = O(n \log n/(1-\lambda_2))$ due to Broder and Karlin [8] from 1989. In all four cases, t_{meet}, t_{coal}, t_{hit}, and t_{cov}, our bounds improve the dependency on $1/(1-\lambda_2)$ (or, equivalently t_{mix}), by almost a square-root (we refer to Theorem 6.5, Theorem 6.6 and Theorem 6.7 for more details).

We also derive a general lower bound on t_{meet} that combines the trivial bound, $1/\|\pi\|_2^2$,with the minimum number of returns (see Theorem 6.1). Although this bound does not directly provide the correct lower bound for binary trees, it forms the basis of a later analysis in .

Finally, we also provide asympotically tight worst-case bounds on t_{meet} and t_{coal}. We show that on any graph the coalescence time must be at least $\Omega(\log n)$ and is no more than $O(n^3)$. For regular (and vertex-transitive) graphs these bounds become $\Omega(n)$ and $O(n^2)$ (See also Table 2 on page 40, which also contains an explanation why these bounds are asymptotically tight.) These two new upper bounds for general and regular graphs complete the picture of worst-case bounds:

Theorem 1.5. *The following hold for graphs of the stated kind*

(i) For any graph G we have $t_{\text{meet}} \in [\Omega(1), O(n^3)]$ and $t_{\text{coal}} \in [\Omega(\log n), O(n^3)]$

(ii) For any regular (and vertex-transitive) graph G we have $t_{\text{meet}}, t_{\text{coal}} \in [\Omega(n), O(n^2)]$

The proof of Theorem 1.5 appears in Section 5.

[4] The exact model is specified in Section 7.5

1.2 Proof Ideas and Technical Contributions

When dealing with processes involving concurrent random walks, a significant challenge is often to get a handle on the behavior of "short" random walks. This challenge appears in several settings, e.g., in the context of cover time of multiple random walks [3, 18], where Efremenko and Reingold [18, Section 6] highlight the difficulty in analyzing the hitting time distribution before its expectation. In the context of concentration inequalities for Markov chains, Lezaud [24, p. 863] points out the requirement to spend at least mixing time steps before taking any samples. Related to that, in property testing, dealing with graphs that are far from expanders has been mentioned as one of the major challenges to test the expansion of the graph by Czumaj and Sohler [17].

In our setting, we also face these generic problems and devise different methods to get a handle on the meeting time distribution before its expectation. Despite our focus being on coalescing and meeting times, several of our approaches can be leveraged to derive new bounds on other random walk quantities such as hitting times or cover times (see Section 6).

Specifically, the key ingredient in our first proof (Theorem 1.1), where we express t_{coal} as a tradeoff between t_{meet} and t_{mix} is the following: a tight bound on the probability p_ℓ that two random walks meet before ℓ time-steps, for ℓ in the range $[t_{\text{mix}}, t_{\text{coal}}]$. Arguing about meeting probabilities of walks that are much shorter than t_{meet} allows us to understand the rate at which the number of *alive* random walks is decreasing.

Optimistically, one may hope that starting with k random walks, as there are $\binom{k}{2}$ possible meeting events, roughly $\binom{k}{2} \cdot p_\ell$ meetings may have occurred after ℓ time-steps. However, the non-independence of these events turns out to be a serious issue and we require a significantly more sophisticated approach to account for the dependencies. We divide the k random walks into disjoint groups G_1 and G_2 (with $|G_1|$ being much smaller than $|G_2|$) and consider only eliminations of random walks in G_2 that collide with a random walk in G_1. The domination of the real process by the group-restricted one is established by introducing a formal concept called *selector functions* (see beginning of Section 3.1 and Figure 2), that allows us for each vertex, each time-step, and each subset of walks appearing at a vertex at a given time-step, to apply a "coalescing" rule (this includes the possibility that none of the walks merge). In this resulting stochastic process, we can expose the random walks of G_1 first and consider meetings with random walks in G_2 (for an illustration, see Figure 2 on page 11.) Conditioning on a specific exposed walk in G_1, the events of the different walks in G_2 meeting this exposed walk are indeed independent. Thus, the problem then reduces to calculating the likelihood of a random walk in G_2 having a 'good trajectory', *i. e.,* one which many random walks in G_1 would meet with large enough probability.

Surprisingly, it suffices to divide trajectories into only two classes (Lemma 3.5). Although, one may expect that a more fine-grained classification of trajectories would result in better bounds, it turns out not to be the case. In fact, the bound that we derive on the coalescing time in Theorem 1.1 is tight, and this is precisely due to the tightness of Lemma 3.5. The tightness is established by the following construction (cf. Figure 1). The graph is designed such that the vast majority of meetings (between any two random walks) occur in a relatively small part of the graph (G_2 in Figure 1). On average, it takes a considerable number of time-steps before random walks actually get to this part of the graph. What this implies is that for relatively short trajectories (of length significantly smaller than t_{meet}), it is quite likely that other random walks will not meet them (cf. Lemma 3.5). There is a bit of a dichotomy here, once a walk reaches G_2 it is likely that many random walks will meet it; however, a random walk not reaching G_2 is unlikely to be met by any other random walk.

Equipped with Theorem 1.1, we can bound $t_{\text{coal}} = \Theta(t_{\text{meet}})$ for all graphs satisfying $t_{\text{meet}}/t_{\text{mix}} \geqslant (\log n)^2$. Therefore, the problem of bounding t_{coal} reduces to bounding t_{meet}.

The starting point in the proof of most of our bounds on t_{meet} is the following simple obser-

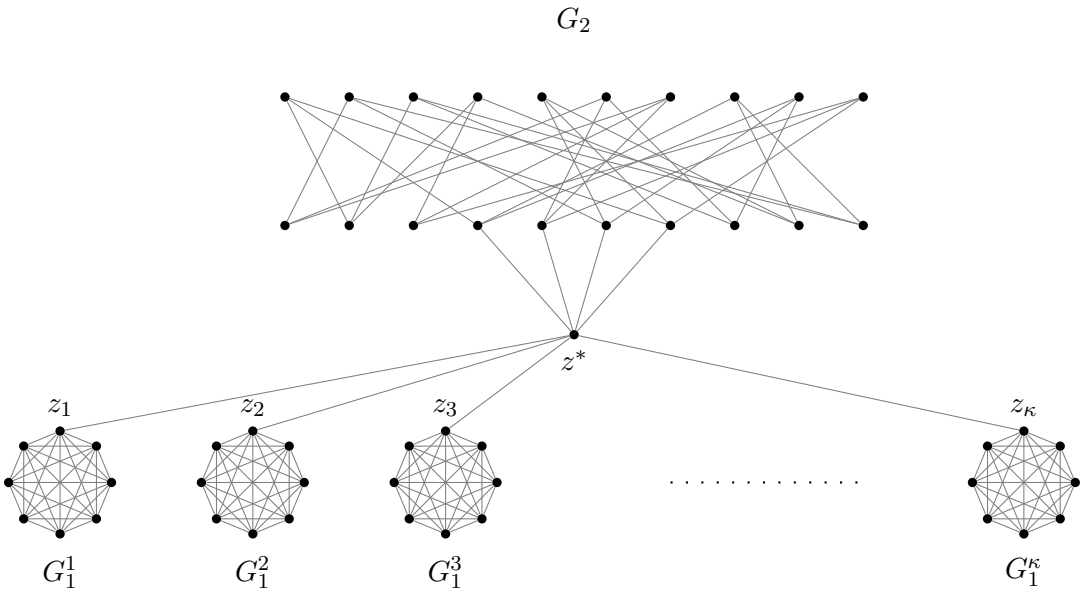

Figure 1: The graph described in Section 3.4 with $t_{\text{coal}} = \Omega(t_{\text{meet}} + \sqrt{\alpha} \cdot \log n \cdot t_{\text{mix}})$.

vation. If we have two random walks $(X_t)_{t \geqslant 0}$ and $(Y_t)_{t \geqslant 0}$, and count the number of collisions $Z := \sum_{t=0}^{\tau} \mathbf{1}_{X_t = Y_t}$ within the first τ time-steps, then

$$\mathbb{P}[\, Z \geqslant 1 \,] = \frac{\mathbb{E}[\, Z \,]}{\mathbb{E}[\, Z \mid Z \geqslant 1 \,]}. \tag{1}$$

If we further assume that both walks start from the stationary distribution, then we have

$$\mathbb{P}[\, Z \geqslant 1 \,] = \frac{\tau \cdot \|\pi\|_2^2}{\mathbb{E}[\, Z \mid Z \geqslant 1 \,]}$$

We should mention that this generic approach is of course not new, an analogous variant of counting visits to a vertex has been used by Cooper and Frieze in several works (e.g., [11]) to derive very accurate bounds on the hitting (and cover time) on various classes of random graphs, or in Barlow et al. [4] to bound the collisions of random walks on infinite graphs. However, using (1), we were able to obtain several improvements to existing bounds on the meeting time, and as a consequence for coalescing time. We believe that our work further highlights the power of this identity.

The crux of (1) is that in order to lower (or upper) bound the probability that the two walks meet, we need to derive a corresponding bound on $\mathbb{E}[\, Z \mid Z \geqslant 1 \,]$, *i. e.,* the number of collisions conditioning on the occurrence of at least one collision.

Our results employ various tools to get a handle on this quantity, but here we mention one that is relatively simple:

$$\mathbb{E}[\, Z \mid Z \geqslant 1 \,] \leqslant \max_{u \in V} \sum_{t=0}^{\tau} \sum_{v \in V} \left(p_{u,v}^t\right)^2. \tag{2}$$

The inner summand $\sum_{v \in V} (p_{u,v}^t)^2$ is the probability that two walks starting from the same vertex u will meet after t steps. Thus, summing over t and conditioning on the first meeting happening (*i. e.,* the condition $Z \geqslant 1$) at some vertex u before time-step τ yields the bound in (2). Despite

the seemingly crude nature of this bound, it can be used to derive new results for $t_{\text{hit}}, t_{\text{meet}}$ and t_{coal} that significantly improve over the state-of-the-art on regular graphs (see Section 6, or the last paragraph in this Section for a summary).

One challenge is that, even if we have grip on $\mathbb{P}[\, Z \geqslant 1\,]$ for different values of τ, it only gives us information about the event that *two* random walks meet. If we extend this to $k \gg 2$ random walks, we irrevocably run into dependency issues. That is, if we have say $k = 3$ random walks, then conditioning on the event that the first random walk does not meet with the second may decrease the probability that the first random walk meets with the third. For instance, such dependencies may arise in graphs where there is a dense core in which random walks are likely to meet, while on the periphery it is very unlikely for the walks to meet (in fact, we will construct one example (Section 3), which despite being an almost-regular graph, maximizes these dependencies). Coping with these inherit dependencies among the k walks is challenging, and we apply a variety of methods that yield different results in the end.

Reduction Techniques

The derivation of our bounds on t_{coal} in terms of t_{hit} are based on two general reduction results, that might be useful in other applications:

Theorem 1.6 (Reduction Results)**.** *The following results hold for any graph G:*

1. *The coalescence process reduces the number walks from n to $O(\log^3 n)$ in $O(t_{\text{hit}})$ steps with probability at least $1 - n^{-1}$. (see Theorem 4.3)*
2. *The coalescence process reduces the number walks from $O(\log^3 n)$ to $(\Delta/d)^{O(1)}$ in $O(t_{\text{hit}})$ steps in expectation, where Δ is the maximum degree and d is the average degree (see Theorem 4.4)*

A basic ingredient are new concentration inequalities, which are derived in Section 4.1. Our concentration inequalities yield sufficiently strong bounds for upper tails of returns (or other, possibly more complex random variables) by a random walk of length t_{hit}, while most of the existing bounds (e.g., [10, 24]) require that the expectation of the random variable is at least as large as t_{mix}. While $t_{\text{mix}} \leqslant t_{\text{hit}}$ in general, the challenging case in our analysis is when $t_{\text{mix}} \approx t_{\text{hit}}$ and in this cases our concentration tools seem to provide stronger upper tails than the existing concentration tools.

Equipped with these concentration results, the proof of Theorem 4.3 is surprisingly simple and rests again on (1). First, by a straightforward bucketing argument on the degree distribution, we show that with high probability, we can find for each random walk $(X_t)_{t\geqslant 0}$ a set S (depending on the trajectory of X_t), so that with high probability, (i) each vertex in $S(i)$ is visited frequently during $O(t_{\text{hit}})$ steps, and (ii) each vertex in $S(i)$ has the same degree up to constants. Conditioning on this, it follows that a second random walk $(Y_t)_{t\geqslant 0}$ will have sufficient collisions in expectations, i.e., $\mathbb{E}[\, Z\,]$ is large enough. To bound $\mathbb{E}[\, Z \mid Z \geqslant 1\,]$, we use the concentration inequalities to establish that with high probability, the trajectory $(X_t)_{t\geqslant 0}$ will be good in the sense that $\mathbb{E}[\, Z \mid (x_0, x_1, \ldots), Z \geqslant 1\,]$ is not too large. Combining these bounds yields $\mathbb{P}[\, Z \geqslant 1\,] = \Omega(1/\log^3 n)$, and a straightforward division into groups G_1 and G_2 of sizes $\Theta(\log^3 n)$ and $n - |G_1|$ shows that all random walks in G_2 can be eliminated in $O(t_{\text{hit}})$ steps.

The proof of the second reduction result (Theorem 4.4) is more involved, although it again revolves around (1). The issue is that we can no longer repeat the simple bucketing argument from Theorem 4.3 about the degree distribution, since the number of buckets may vastly exceed the number of walks. Furthermore, we may no longer obtain "w.h.p."-bounds on the probability for certain good events. For all these reasons, a refined approach is needed.

Our analysis allocates small phases of length $O(t_{\text{hit}}/\kappa)$ in order to halve the number of random walks, where $k = \kappa^{100}$ is the number of walks at the beginning of the phase. The first step is to show that starting from any vertex, there exists a large set of vertices, so that each vertex is visited the "right" amount of time, but also that it was not too unexpected to visit that vertex. The latter condition is quite subtle, but it allows us to arrange a proper scheduling of the walks to show that, regardless of which vertices the random walk i decides to visit in that set, there are enough walks that are able to reach these vertices by then. In other words, it rules out the possibility that, despite two random walks visiting the same set of vertices, they never collide (for an illustration, see Figure 4 on page 36). Using our concentration bounds with a careful choice of the slackness parameters in terms of κ, the above approach can eventually be shown to reduce the number of random walks k by a constant within $O(t_{\text{hit}}/\kappa)$ steps. Repeating this iteratively yields the bound $O(t_{\text{hit}})$.

Bounds on t_{hit} and Worst-Case Bounds

With the two reduction results, Theorem 1.3 follows immediately. Furthermore, the aforementioned results can be also used to derive worst-case upper and lower bounds on meeting and coalescing time on general and regular graphs that are tight up to constant factors. Some of these were known, or follow directly from existing results, the others are, to the best of our knowledge, novel.

We proceed by establishing that $t_{\text{coal}} = O(n^3)$ on all graphs. The proof of $t_{\text{coal}} = O(n^3)$ follows by first applying both reductions (Theorem 4.3 and Theorem 4.4) to reduce the number of walks from n to $(\Delta/d)^{O(1)} \leqslant (n^2/|E|)^{O(1)}$ in $O(t_{\text{hit}})$. Finally, combining the bound $t_{\text{meet}} \leqslant 4t_{\text{hit}} = O(n \cdot |E|)$ together with $t_{\text{coal}}(S_0) = O(t_{\text{meet}} \cdot \log(|S_0|))$ (Proposition 3.3) for any set of start vertices S_0, yields that after additional

$$O(t_{\text{meet}} \cdot \log(|S_0|)) = O(n \cdot |E| \cdot \log(n^2/|E|)) = O(n^3)$$

steps the coalescing terminates. The fact that this is tight can be easily verified by considering the Barbell graph [5].

For regular graphs, the same argument as before shows that $t_{\text{coal}} = O(n^2)$, and this is matched by the cycle, for instance. The proof of the other results are straightforward, and we refer the reader to Section 5.3.

Bounds on t_{meet} and Other Results

In Section 6, we derive several bounds on t_{meet}. These bounds are derived more directly by (1) and/or (2), and involve other quantities such as $\|\pi\|_2^2$ or the eigenvalue gap $1 - \lambda_2$. One important technical contribution is to combine routine spectral methods involving the spectral representation and fundamental matrices that have been used in previous works, *e. g.,* Cooper et al. [14] with some short-time bounds on the t-th step probabilities. This allows us to improve several bounds, not only on t_{meet} and t_{coal} but also t_{hit} and t_{cov}, by significantly reducing the dependency on the spectral gap or mixing time by almost a square root factor. As a corollary, we also derive a new bound on the cover time for regular graphs that improves over the best known bound by Broder and Karlin [8] from 1989.

[5] This n-vertex graph is constructed by taking two cliques of size $n/4$ each, and connecting them through a path of length $n/2$.

Concrete Topologies

Finally, in Section 7 we apply the derived upper and lower bounds on t_{meet} and t_{coal} on various fundamental topologies including grids, expanders and hypercubes. In most cases, these results follow immediately from the general bounds by plugging in corresponding values for $\|\pi\|_2^2$, t_{hit} or t_{mix}. One exception is the binary tree, for which it seems surprisingly non-trivial to derive a lower bound of $t_{\text{meet}} = \Omega(n \log n)$. Here again we use a refinement of (1) that restricts the vertices to leaf-nodes u, for which $\sum_{t=1}^{t_{\text{mix}}} (\sum_{v \in V} p_{u,v}^t)^2 = \Omega(\log n)$. The matching upper bound $t_{\text{meet}} = O(n \log n)$ follows from $t_{\text{coal}} = O(t_{\text{hit}})$.

Of particular interest might be the analysis of "real-world" graph models given in Section 7.5. There we show how to utilize our bounds from earlier sections to establish $t_{\text{coal}} = \Theta(t_{\text{meet}})$ on two random graph models, leading to bounds on t_{coal} that are sublinear in the number of vertices.

1.3 Discussion and Future Work

In this work we derived several novel bounds on t_{coal}. Our first main result implies that a gap of just $\Omega(\log^2 n)$ between t_{mix} and t_{meet} is sufficient to have $t_{\text{coal}} = \Theta(t_{\text{meet}})$. We also proved that this result is essentially tight. Further, we derived several new bounds on t_{coal} based on t_{hit}. For almost-regular-graphs, our new result implies the following hierarchy for the discrete-time setting

$$t_{\text{meet}} \leqslant t_{\text{coal}} = O(t_{\text{hit}}),$$

which refines the already known result $t_{\text{meet}} = O(t_{\text{hit}})$. Finally, we also determined tight worst-case lower and upper bound for t_{coal}.

For future work, an obvious problem is to extend the $t_{\text{coal}} = O(t_{\text{hit}})$ result to all graphs (so far, we only know $t_{\text{coal}} = O(t_{\text{hit}} \cdot \log \log n)$). Even more ambitious would be to try to link the continuous-time variant to the discrete-time process, as this would immediately resolve the $t_{\text{coal}} = O(t_{\text{hit}})$ problem. A different direction may be to further explore lower bounds on t_{meet}; in this work we only derived one lower bound on t_{meet} in Theorem 6.1.

2 Notation and Preliminaries

Throughout the paper, let $G = (V, E)$ denote an undirected, connected graph with $|V| = n$ and $|E| = m$. For a node $u \in V$, $\deg(u)$ denotes the degree of u and $N(u) = \{v \colon (u, v) \in E\}$ the *neighborhood* of u. By Δ and $d = \frac{1}{n} \sum_{u \in V} \deg(u)$, we denote the maximum and average degree, respectively. Unless stated otherwise, all random walks are assumed to be discrete-time (indexed by natural numbers) and lazy, *i. e.,* if P denotes the $n \times n$ transition matrix of the random walk, $p_{u,u} = 1/2$, $p_{u,v} = 1/(2 \deg(u))$ for any edge $(u, v) \in E$ and $p_{u,v} = 0$ otherwise. We define $p_{u,v}^t$ to be the probability that a random walk starting at $u \in V$ is at node $v \in V$ at time $t \in \mathbb{N}$. Furthermore, let $p_{u,\cdot}^t$ be the probability distribution of the random walk after t steps starting at u. By π we denote the *stationary distribution*, which satisifes $\pi(u) = \deg(u)/(2m)$ for all $u \in V$.

Let $d(t) = \max_u \|p_{u,\cdot}^t - \pi\|_{\text{TV}}$ and $\bar{d}(t) = \max_{u,v} \|p_{u,\cdot}^t - p_{v,\cdot}^t\|_{\text{TV}}$, where $\|\cdot\|_{\text{TV}}$ denotes the total variation distance. Following Aldous and Fill [2], we define the *mixing time* to be $t_{\text{mix}}(\varepsilon) = \min_t \{t \geqslant 0 : \bar{d}(t) \leqslant \varepsilon\}$ and for convenience we will write $t_{\text{mix}} = t_{\text{mix}}(1/e)$. We define separation from stationarity to be $s(t) = \min\{\varepsilon \colon p_{u,v}^t \geqslant (1-\varepsilon)\pi(v) \text{ for all } u, v \in V\}$. Then $s(.)$ is submultiplicative, so in particular, non-increasing [2], and we can define the *separation threshold time* $t_{\text{sep}} = \min_t \{t \geqslant 0 : s(t) \leqslant e^{-1}\}$ and, by [2, Lemma 4.11], $t_{\text{sep}} \leqslant 4t_{\text{mix}}$. The *hitting time* $t_{\text{hit}}(u, v)$ of any pair of nodes $u, v \in V$ is the expected time required for a random walk starting at u to hit v. The hitting time $t_{\text{hit}} = \max_{u,v} t_{\text{hit}}(u, v)$ is the maximum over all such pairs. For $A \subseteq V$, we use

$t_{\text{hit}}(u, A)$, to denote the expected time required for a random walk starting to u to hit some node in the set A. Furthermore, we define $t_{\text{hit}}(\pi, u) = \sum_{v \in V} t_{\text{hit}}(v, u) \cdot \pi(v)$. Furthermore, we define $t_{\text{avg-hit}} = \sum_{u,v \in V} \pi(u) \cdot \pi(v) \cdot t_{\text{hit}}(u, v)$.

Let $t_{\text{meet}}(u, v)$ denote the expected time when two random walks starting at u and v arrive at the same node for the first time, and we write t^{π}_{meet} for the expected meeting time of two random walks starting at two independent samples from the stationary distribution. Finally, let $t_{\text{meet}} = \max_{u,v} t_{\text{meet}}(u, v)$ denoted the worst-case expected meeting time.

We define the coalescence process as a stochastic process as follows: Let $S_0 \subseteq V$ and for all $v \in S_{t-1}$ let

$$Y^t_v = \begin{cases} u \in N(v) & \text{w.p. } \frac{1}{2|N(v)|} \\ v & \text{w.p. } \frac{1}{2} \end{cases}$$

The set of *active* nodes in step t is given by $S_t = \{Y^t_v \mid v \in S_{t-1}\}$. The process satisfies the Markov property, *i. e.*,

$$\mathbb{P}[\, S_t \mid \mathcal{F}_t \,] = \mathbb{P}[\, S_t \mid S_{t-1} \,], \tag{3}$$

where $\mathcal{F}_t$ is the filtration up to time t. Finally, we define the *time of coalescence* as $T(S_0) = \min\{t \geqslant 0 \mid |S_t| = 1\}$. Let $t_{\text{coal}}(S_0)$ be defined as $\mathbb{E}[T_{S_0}]$; we use the notation t_{coal} when $S_0 = V$. While this definition is mathematically more convenient, the reader can verify that it is indeed equivalent to running multiple random walks starting at nodes in S_0 and merging them two or more meet at a node.

Throughout this paper, the expression w.h.p. (*with high probability*) means with probability at least $1 - n^{-\Omega(1)}$ and the expression w.c.p. (*with constant probability*) means with probability > 0. We use $\log n$ for the natural logarithm. Appendix A contains known results about Markov Chains that we frequently use in our proofs.

3 Bounding t_{coal} for large $t_{\text{meet}}/t_{\text{mix}}$

In this section we prove Theorem 1.1, our main theorem. We refer the reader to Section 1.2 for a high-level description of the proof ideas.

3.1 Stochastic Process

In order to prove our main result, it is helpful to consider a more general stochastic process involving multiple independent random walks. In the more general process, whenever several random walks arrive at the same node at the same time a subset of them (rather than just one) may survive, while the remaining are merged with one of the surviving walks. To identify the random walks, we assume that each walk has a natural number (in $\mathbb{N}$) as an identifier. In order to define this process formally, we introduce some additional notation and definitions; then we state and prove some auxiliary lemmas. A similar but slightly less general concept was introduced in [29].

Definition 3.1. *For a graph $G = (V, E)$, a set of functions $(\sigma_v)_{v \in V}$, where for each $v \in V$, $\sigma_v : 2^{\mathbb{N}} \to 2^{\mathbb{N}}$, is called a* selector function family *if the following are satisfied: (i) $\forall v, S$ we have $\sigma_v(S) \subseteq S$, and (ii) $\forall v, S \neq \emptyset$ we have $\sigma_v(S) \neq \emptyset$. For selector function families $(\sigma_v), (\sigma'_v)$ we denote $(\sigma_v) \prec (\sigma'_v)$ if $\forall v, S, \sigma_v(S) \subseteq \sigma'_v(S)$.*

As mentioned before, we assume that every random walk r has a unique identifier $\mathsf{id}(r) \in \mathbb{N}$. The selector function at a *node* determines which of the random walks that arrive at this node at the same time survive. We will make use of the generality of the definition later in Section 5. In

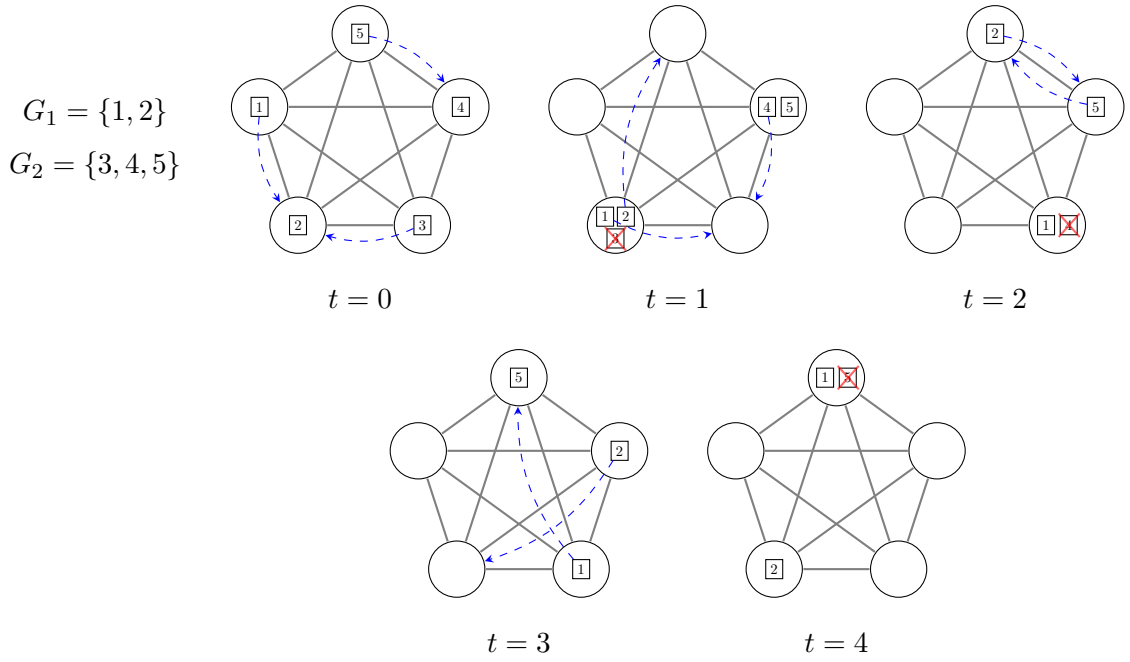

Figure 2: Illustration of the selector function $\sigma^*(S) = S \setminus G_2$ if $S \cap G_1 \neq \emptyset$ and $\sigma^*(S) = S$ otherwise. In words, if some random walks with ids in G_1 and some with ids in G_2 are present on the same node at the same time, only the ones with ids in G_1 survive. However, if all the random walks have ids in only one of G_1 or G_2, then all of them survive.

this section, we will only make use of selector function families where σ_v are the same for all v and simply call this a selector function σ; furthermore, we will assume that for $S \neq \emptyset$ it is the case that $\min\{i \in S\} \in \sigma(S)$, which ensures that $\sigma(S) \neq \emptyset$. The choice of the walk with smallest id surviving in the definition of the selector function is arbitrary; we only require that at least one walk survives. We denote by P the stochastic process that keeps track of positions of random walks along with their ids. Formally, let P denote the following process:

1. At time 0, $S_0 = \{(u_r, \mathsf{id}(r))\}$, where u_r is the starting node of random walk r and $\mathsf{id}(r)$ is its identifier.

2. At time t, several random walks may arrive at the same node. The process P allows some subset of them to survive, while the rest 'coalesce' with one of the surviving walks. Formally, S_{t+1} is defined using S_t as follows. Let σ be a selector function and define the target of the random walk with id i which on node v to be

$$Y_{v,i} = \begin{cases} (u, i) \text{ where } u \in N(v) & \text{w.p. } \frac{1}{2|N(v)|} \\ (v, i) & \text{w.p. } \frac{1}{2}, \end{cases}$$

for $i \in \mathbb{N}$ and $v \in V$. Let $R_t = \{(Y_{v,i}, i) \mid (v, i) \in S_t\}$ be the set of targets (before merging happens). Then, merging happens w.r.t. to the selector as follows

$$S_{t+1} = \bigcup_{v \in V} \{(v, j) \mid j \in \sigma(\{i \mid (v, i) \in R_t\})\}.$$

We consider a special process P^c based on P: the selector function is defined as $\sigma^c(S) = \{\min_{i \in S} i\}$ and the starting nodes u_r in S_0 are all distinct. In process P^c, only one of several walks arriving at the same node survives and by convention the one having the smallest id is chosen. Let $(S_t^c)_{t=0}^{\infty}$ denote the stochastic process P^c. If $\bar{S}_t^c = \{v \mid (v,i) \in S_t^c\}$, then $(\bar{S}_t^c)_{t=0}^{\infty}$ is a coalescence process as defined in Section 2.

Observation 3.2. *Consider two copies of process P, denoted by P and P' with starting sets S_0 and S_0' and selector functions σ and σ' respectively. If $S_0 \subseteq S_0'$ and $\sigma \prec \sigma'$, there exists a coupling such that for all t, $S_t \subseteq S_t'$.*

It is well-known that starting with k random walks, the coalescence time is bounded by $O(t_{\text{meet}} \log k)$, this can be deduced from the proof presented in [22]. For the sake of completeness, we give a proof.

Proposition 3.3. $t_{\text{coal}}(S_0) = O(t_{\text{meet}} \log |S_0|)$.

Proof. Let P be the process defined in Section 3.1. Let S_t be set of coalescing random walks at an arbitrary time-step t.

In the following we show the slightly stronger claim that expected time to reduce the number of random walks by a constant factor is of order $O(t_{\text{meet}})$. Formally, we fix an arbitrary time-step t_0. Let $\tau = \min\{t \geqslant t_0 \colon |S_t| \leqslant 11/12 \cdot |S_{t_0}|, |S_{t_0}| \geqslant 10\}$ be the first time-step the number of coalescing random walks reduces by a factor of $11/12$, then we show $E[\tau] = O(t_{\text{meet}})$. Iterating the argument $O(\log |S_0|)$ times implies that the expected time it takes to reduce to 10 random walks is $O(t_{\text{meet}} \log |S_0|)$. Note that the expected time to reduce from 10 random walks to 1 is bounded by $O(t_{\text{meet}})$. Hence, the claim $t_{\text{coal}}(S_0) = O(t_{\text{meet}} \log |S_0|)$ follows.

It remains to show that the expected number of rounds it takes to reduce the number of random walks by a factor of $11/12$ is indeed $O(t_{\text{meet}})$. We divide time into blocks of length $T = 2t_{\text{mix}}$. Let $X_\ell = |S_{t_0+\ell \cdot T}|$ denote the number of random walks alive at the beginning of block ℓ. For convenience we assume the the random walks in S_{t_0} have distinct ids in $[1, |S_{t_0}|]$. We construct a maximal matching between random walks and only allow coalescence between matched random walks. Formally, let $M = \{(i, i+1) \colon i \bmod 2 = 1,\ \exists u, v \in V s.t.\ (u,i) \in S_{t_0+\ell \cdot T},\ (v, i+1) \in S_{t_0+\ell \cdot T}\}$. Let A and B the corresponding sets, *i. e.*, $A = \{i : (i,j) \in M\}$ and $B = \{j : (i,j) \in M\}$. Let σ be the selector function defined the following way: For $S \subseteq 2^{\mathbb{N}}$ define

$$\sigma(S) = S \setminus \{(u, i+1) \colon u \in V, (u,i) \in S, (u, i+1) \in S\}.$$

Observe that $|A| \geqslant X_0/2 - 1$ and $X_\ell \geqslant 11/12 \cdot X_0$ for $\tau + \ell \cdot T \leqslant \tau$.

Furthermore, after any period of length T, two random walks starting from arbitrary positions meet, by Markov inequality, w.p. at least $1/2$. Hence, we have for the process P with the above defined σ selector function that

$$\begin{aligned}\mathbb{E}\left[X_\ell \mid \mathcal{F}_{t_0+(\ell-1)\cdot T}, \tau > t_0 + (\ell-1)\cdot T \right] &\leqslant X_\ell - \frac{|A \cap X_{\ell-1}|}{2} \leqslant X_\ell - \frac{|A| - X_0/12}{2} \\ &\leqslant X_\ell - \frac{X_0/2 - 1 - X_0/12}{2} \leqslant 11/12 \cdot X_0\end{aligned}$$

We can therefore apply Lemma A.7 with parameters $g = 11/12 \cdot X_0$ and $\beta = 11/12$ we obtain that $t_{\text{coal}}(S_0) \leqslant O(T) = O(t_{\text{meet}})$, which completes the proof. ∎

3.2 Meeting Time Distribution Prior to t_{meet}

Let $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$ be independent random walks starting at arbitrary positions. For τ a multiple of t_{mix}, the following lemma gives a lower bound on the probability of intersection of the two random walks in τ steps.

Lemma 3.4. *Let $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$ be two independent random walks starting at arbitrary positions. Let* $\text{intersect}(X_t, Y_t, \tau)$ *be the event that there exists $0 \leqslant s \leqslant \tau$, such that $X_s = Y_s$. Then*

$$\mathbb{P}\left[\,\text{intersect}(X_t, Y_t, 5t_{\text{mix}})\,\right] \geqslant \frac{1}{16\alpha},$$

where $\alpha = t_{\text{meet}}/t_{\text{mix}}$. Furthermore, for any $1 \leqslant b \leqslant \frac{e-1}{e}\cdot\alpha$, there exists a constant c, such that

$$\mathbb{P}\left[\,\text{intersect}(X_t, Y_t, cbt_{\text{mix}})\,\right] \geqslant \frac{b}{\alpha},$$

Proof. First, let $(\widetilde{X}_t)_{t\geqslant 0}$ and $(\widetilde{Y}_t)_{t\geqslant 0}$ be two random walks that from two independent samples from the stationary distribution and are run for $\ell := 2\lceil\alpha\rceil t_{\text{mix}}$ steps. Notice that $\ell \geqslant 2t_{\text{meet}}$, and hence, by Markov's inequality,

$$\mathbb{P}\left[\,\text{intersect}(\widetilde{X}_t, \widetilde{Y}_t, \ell)\,\right] \geqslant \frac{1}{2}. \tag{4}$$

Furthermore, if we divide the interval $[1, \ell]$ into $2\lceil\alpha\rceil$ consecutive sections of length t_{mix} each, the probability for a collision in each of these section is identical and therefore the union bound implies

$$\mathbb{P}\left[\,\text{intersect}(\widetilde{X}_t, \widetilde{Y}_t, \ell)\,\right] \leqslant 2\lceil\alpha\rceil \cdot \mathbb{P}\left[\,\text{intersect}(\widetilde{X}_t, \widetilde{Y}_t, t_{\text{mix}})\,\right], \tag{5}$$

and hence combining equation (4) and (5) yields

$$\mathbb{P}\left[\,\text{intersect}(\widetilde{X}_t, \widetilde{Y}_t, t_{\text{mix}})\,\right] \geqslant \frac{1}{4\lceil\alpha\rceil}.$$

Consider now two independent random walks $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$ starting at arbitrary positions. By applying Lemma A.4 to both walks, with probability at least $(1-e^{-1})^2$ both $X_{4t_{\text{mix}}}$ and $Y_{4t_{\text{mix}}}$ are drawn independently from the stationary distribution. Therefore,

$$\mathbb{P}\left[\,\text{intersect}(X_t, Y_t, 5t_{\text{mix}})\,\right] \geqslant (1-e^{-1})^2 \cdot \mathbb{P}\left[\,\text{intersect}(\widetilde{X}_t, \widetilde{Y}_t, t_{\text{mix}})\,\right] \geqslant (1-e^{-1})^2 \cdot \frac{1}{4\lceil\alpha\rceil}.$$

Observing that for any $\alpha \geqslant 1$, the RHS above expression is greater than $1/(16\alpha)$ completes the proof of the first part.

For the second part, we consider k blocks of length $5t_{\text{mix}}$. Using the fact that $(1-x)^k \leqslant 1-(1-e^{-1})xk$ for $0 \leqslant x < 1$, $k \geqslant 0$ and $xk < 1$, we derive that the probability of intersecting after $k = \left\lceil 16b/(1-e^{-1})\right\rceil$ blocks is at least $1-(1-1/(16\alpha))^k \geqslant (1-e^{-1})k/(16\alpha) = b/\alpha$ yielding the second part. ∎

At the heart of the proof lies the following lemma that analyses the marginal distribution of the meeting time distribution. That is, we only expose the first random walk $(X_t)_{t=0}^{\tau}$, and look at how this affects the probability of meeting. In essence we show that at least one of the two orthogonal cases hold. In Case 1 (corresponding to set C_1), there is at least a modest probability that after exposing (X_t), (Y_t) will intersect with significant probability. Otherwise, in Case 2, there is a significant probability that after exposing (X_t), (Y_t) will intersect with at least a modest probability.

Lemma 3.5. *Fix $\tau \in \mathbb{N}$. Let $(X_t)_{t=0}^{\tau}$ and $(Y_t)_{t=0}^{\tau}$ be independent random walks, where the starting nodes X_0 and Y_0 are drawn independently from the stationary distribution π, and the walks are run for τ steps. Let $p = \mathbb{P}[\,\text{intersect}(X_t, Y_t, \tau)\,]$. Let $(x_0, \ldots, x_\tau)$ denote a walk in G (including possible self-loops). Let*

$$C_1 = \{(x_0, \ldots, x_\tau) \mid \mathbb{P}[\,\exists 0 \leqslant s \leqslant \tau, Y_s = x_s\,] \geqslant \sqrt{p}\}$$
$$C_2 = \{(x_0, \ldots, x_\tau) \mid \mathbb{P}[\,\exists 0 \leqslant s \leqslant \tau, Y_s = x_s\,] \geqslant \frac{p}{3}\}$$

Then either $\mathbb{P}[\,(X_t)_{t=0}^{\tau} \in C_1\,] \geqslant \frac{p}{3}$ or $\mathbb{P}[\,(X_t)_{t=0}^{\tau} \in C_2\,] \geqslant \frac{\sqrt{p}}{3}$.

While the actual lower bounds on the probabilities appear rather crude, it turns out that the "significant probability" $\sqrt{p}/3$ is best possible, as we demonstrate in our lower bound construction later. Remarkably, the fact that the "modest probability" is only $p/3$ and much smaller than $\sqrt{p}/3$ does not affect the tightness of our bound, since in Claim 3.6, we can make up for this gap in both cases through a simple amplification argument over the unexposed random walks.

Proof. Let us suppose that $\mathbb{P}[\,(X_t)_{t=0}^{\tau} \in C_1\,] < \frac{p}{3}$. We show that this implies $\mathbb{P}[\,(X_t)_{t=0}^{\tau} \in C_2\,] \geqslant \frac{\sqrt{p}}{3}$. Assume for the sake of contradiction $\mathbb{P}[\,(X_t)_{t=0}^{\tau} \in C_2\,] < \frac{\sqrt{p}}{3}$. We have

$$\begin{aligned} p &\leqslant \mathbb{P}[\,\text{intersect}(X_t, Y_t, \tau)\,] \\ &\leqslant \mathbb{P}[\,(X_t)_{t=0}^{\tau} \in C_1\,] \cdot 1 + \mathbb{P}[\,(X_t)_{t=0}^{\tau} \in (C_2 \setminus C_1)\,] \cdot \sqrt{p} + \mathbb{P}[\,(X_t)_{t=0}^{\tau} \notin C_2\,] \cdot \frac{p}{3} \\ &< p/3 + \sqrt{p}/3 \cdot \sqrt{p} + p/3 \leqslant p, \end{aligned}$$

a contradiction. This completes the proof. □

3.3 Upper Bound - Proof of Theorem 1.1

We commence by defining a specific instance of the process P defined in Section 3.1. This allows us to establish Claim 3.6 providing us with the following tradeoff. For a given period ℓ of length at least t_{mix} we obtain a bound on the required time to reduce the number of random walks by an arbitrary factor. The proof relies heavily on Lemma 3.5 which divides the walks of G_1 into two groups allowing us to expose the walks of G_1 and to calculate the probability of the walks of G_2 to intersect with them. These probabilities are derived from the time-probability tradeoff presented in Lemma 3.4. We then use Claim 3.6 to derive a bound on the number of time steps it takes to reduce the number of walks to $\lceil 2\alpha \rceil$, where $\alpha = t_{\text{meet}}/t_{\text{mix}}$ (Lemma 3.7). From there on we employ Claim 3.6 to reduce the number of walks to 1 in $O(t_{\text{meet}})$ time steps. Melding both phases together together yields the bound of Theorem 1.1.

We now define a process $P^*(S_0, k)$ with $k < |S_0|$, which is a special case of the process P defined in Section 3.1, where $P^*(S_0, k)$ is defined as follows:

- Let $|S_0| = k'$; there are k' random walks with ids $1, \ldots, k'$ and starting nodes $v_1, \ldots, v_{k'}$. Thus, $S_0 = \{(v_i, i) \mid 1 \leqslant i \leqslant k'\}$.
- Let $G_1 = \{1, \ldots, k\}$ and $G_2 = \{k+1, \ldots, k'\}$. We define the selector function as follows:

$$\sigma^*(S) = \begin{cases} S \setminus G_2 & \text{if } S \cap G_1 \neq \emptyset \\ S & \text{otherwise} \end{cases}$$

  In words, if some random walks with ids in G_1 and some with ids in G_2 are present on the same node at the same time, only the ones with ids in G_1 survive. However, if all the random walks have ids in only one of G_1 or G_2, then all of them survive.

We define

$$\mathsf{IDs}(S_t) := \{\mathsf{id}(r) \mid (u_r, \mathsf{id}(r)) \in S_t\}, t \in \mathbb{N}$$

Claim 3.6. *Let $\tau \in \mathbb{N}$, let $(X_t)_{t=0}^{\tau}$ and $(Y_t)_{t=0}^{\tau}$ be independent random walks run for τ steps, with X_0 and Y_0 drawn independently from π. Let $p_\tau \leqslant \mathbb{P}[\,\text{intersect}(X_t, Y_t, \tau)\,]$ be a lower bound on the probability of the intersection of the two walks during the τ steps. Consider an instantiation of $P^*(S_0, k)$. Suppose that $k \geqslant \frac{3}{(1-e^{-1})\cdot p_\tau}$. For some $1 \leqslant g \leqslant |S_0| - k$, define the stopping condition $T_g = \min\{t \geqslant 0 \mid |\mathsf{IDs}(S_t) \cap G_2| \leqslant g\}$. Then the expected stopping time is given by*

$$\mathbb{E}[\,T_g\,] = O\left((4t_{\text{mix}} + \tau) \cdot \sqrt{\frac{1}{p_\tau}} \cdot (\log|G_2| - \log g)\right).$$

Proof. We will consider the process in *blocks* each consisting of $4t_{\text{mix}}+\tau$ time-steps. For convenience in the proof, we'll restart counting time-steps from 0 at the beginning of each block; we keep track of the total number of time-steps by counting the number of blocks. Let C_1 and C_2 be as defined in Lemma 3.5. Then we perform a case analysis by considering the two possible outcomes described in Lemma 3.5 separately. We define $Z_j = |\mathsf{IDs}(S_{j\cdot(4t_{\text{mix}}+\tau)}) \cap G_2|$, *i. e.,* the number of walks remaining in G_2 after j blocks of time have passed, for $j \geqslant 1$, we show that there exists a constant $c > 0$ such that,

$$\mathbb{E}[\,Z_j \mid \mathcal{F}_{j-1}\,] \leqslant Z_{j-1} \cdot (1 - c\sqrt{p_\tau}).$$

By using Lemma A.7, we get $\mathbb{E}[\,T_g\,] = O\left((4t_{\text{mix}} + \tau) \cdot \frac{1}{\sqrt{p_\tau}} \cdot (\log|G_2| - \log g)\right)$ (the factor $(4t_{\text{mix}} + \tau)$ appears as the size of the block). In the remainder we show that we have indeed $\mathbb{E}[\,Z_j \mid \mathcal{F}_{j-1}\,] \leqslant Z_{j-1} \cdot \left(1 - c\sqrt{p_\tau}\right)$.

Case 1. $\mathbb{P}[\,(X_t)_{t=0}^{\tau} \in C_1\,] \geqslant \frac{p_\tau}{3}$:

Consider any random walk r in G_1 at the beginning of a *block*. Using Lemma A.4, after $4t_{\text{mix}}$ steps we can couple the state of the random walk with a node drawn from π with probability at least $(1-e^{-1})$. Furthermore, conditioned on this coupling, the portion of the random walk between time-steps $4t_{\text{mix}}$ and $4t_{\text{mix}}+\tau$ of the walk is in class C_1 with probability at least $\frac{p_\tau}{3}$. Since $k \geqslant \frac{3}{p_\tau\cdot(1-e^{-1})}$, w.c.p., in any block, there exists a walk in G_1 that has the portion between time-steps $4t_{\text{mix}}$ and $4t_{\text{mix}} + \tau$ in C_1.

Fix a block and condition on the event that there is a walk in G_1, denoted by r_1, whose portion between time-steps $4t_{\text{mix}}$ and $4t_{\text{mix}} + \tau$ is in C_1. Consider any walk in G_2, denoted by r_2, at the beginning of the block. We want to argue that this walk r_2 has a reasonable probability of intersecting some walk in G_1 in this block of time-steps. First, consider (the possibly hypothetical continuation of r_2) walk r_2' for the entire length of the block. The reason for this is that if r_2 and some walk from G_1 are at the same node at the same time sometime in the block, r_2 will be destroyed in the process $P^*(S_0, k)$; however, we can consider its hypothetical extension to the entire length of the block. Using Lemma A.4 the state of the walk r_2' at time-step $4t_{\text{mix}}$ can be coupled with a node drawn from π with probability at least $1-e^{-1}$. Then conditioned on successful coupling, the probability that r_2' and r_1 collide during time-steps $4t_{\text{mix}}$ and $4t_{\text{mix}}+\tau$ is at least $\sqrt{p_\tau}$ (by definition of C_1 in Lemma 3.5). Thus, the probability that r_2 hits some walk in G_1 is at least $c\sqrt{p_\tau}$ for some constant $c > 0$.

Case 2. $\mathbb{P}[\,(X_t)_{t=0}^{\tau} \in C_2\,] \geqslant \frac{\sqrt{p_\tau}}{3}$:

With slight abuse of notation, consider a walk in G_2, denoted by r_2, at the beginning of a block; as in the previous case, we will consider a possibly hypothetical continuation r_2' of r_2. Using Lemma A.4 we can couple the state of r_2' at time-step $4t_{\text{mix}}$ with a node drawn from π with probability at

least $1 - e^{-1}$. Furthermore, conditioned on the successful coupling, with probability at least $\frac{\sqrt{p_\tau}}{3}$ the trajectory of r'_2 between the time-steps $4t_{\text{mix}}$ to $4t_{\text{mix}} + \tau$ is in C_2. Thus, with probability at least $(1 - e^{-1})\frac{\sqrt{p_\tau}}{3}$, r'_2 has a trajectory between time-steps $4t_{\text{mix}}$ and $4t_{\text{mix}} + \tau$ that lies in C_2. Now consider any random walk $r_1 \in G_1$ at the beginning of the block. Again, using Lemma A.4 with probability at least $1 - e^{-1}$, we can couple the state of the random walk at time $4t_{\text{mix}}$ with a node drawn from π. Conditioned on this between time-steps $4t_{\text{mix}}$ to $4t_{\text{mix}} + \tau$, this random walk hits any trajectory whose portion between time-steps $4t_{\text{mix}}$ to $4t_{\text{mix}} + \tau$ lies in C_2 with probability at least $p_\tau/3$ (by definition of C_2 in Lemma 3.5). Since $k = |G_1| \geqslant \frac{3}{(1-e^{-1})\cdot p_\tau}$, with at least constant probability there is some walk in G_1 that intersects any fixed trajectory whose portion between time-steps $4t_{\text{mix}}$ to $4t_{\text{mix}} + \tau$ lies in C_2. Since the random walks in G_1 are independent, by the definition of the selector function, we have that any walk in G_2 is destroyed by the end of the block with probability at least $c\sqrt{p_\tau}$ for some constant $c > 0$. Thus, we get

$$\mathbb{E}\left[\, Z_j \mid \mathcal{F}_{j-1} \,\right] \leqslant Z_{j-1} \cdot (1 - c\sqrt{p_\tau}) \,.$$

□

Lemma 3.7. *Consider the coalescence process starting with set S_0 and let $\alpha = t_{\text{meet}}/t_{\text{mix}}$. Let $T_1 = \min\{t \geqslant 0 \mid |S_t| \leqslant 2\lceil\alpha\rceil\}$. Then $\mathbb{E}[T_1] = O(t_{\text{mix}} \cdot \sqrt{\alpha} \cdot \log|S_0|)$.*

Proof. We consider the process P^c (defined in Section 3.1), which is identical to the coalescence process, but in addition also keeps track of ids of random walks. What is particular about P^c is that it uses a selector function that allows only the walk with the smallest id to survive. We assume that the ids are from the set $\{1, 2, \ldots, |S_0|\}$. Thus, if $S_0 = \{v_1, \ldots, v_{|S_0|}\}$, $S_0^c = \{(v_1, 1), \ldots, (v_{|S_0|}, |S_0|)\}$. We consider the process $P^*(S'_0, k)$, with $S'_0 = S_0^c$ and $k = \lceil\alpha\rceil$. Let T_1^* be the stopping time defined by $|\mathsf{IDs}(S'_t) \cap G_2| \leqslant \alpha$ for the process $P^*(S'_0, k)$. By definition of P^*, it follows that $\sigma^c \prec \sigma^*$ and thence using Observation 3.2, it follows that T_1^* stochastically dominates T_1. Thus, it suffices to bound $\mathbb{E}[T_1^*]$. We apply Lemma 3.4 with $b = 8$ and derive that for some suitable constant c,

$$p = \mathbb{P}\left[\,\text{intersect}(X_{t\geqslant 0}, Y_{t\geqslant 0}, 8ct_{\text{mix}})\,\right] \geqslant \frac{8}{\alpha},$$

Thus, we have

$$\frac{3}{(1-e^{-1})\cdot p} \leqslant \frac{3}{\frac{1}{2}\cdot p} \leqslant \alpha \leqslant k$$

Applying Claim 3.6 with $g = \alpha$, $\tau = 8ct_{\text{mix}}$ (where c is a constant as given by Lemma 3.4), $p_\tau = 8/\alpha$, and observing that $k \geqslant \frac{3}{(1-e^{-1})\cdot p_\tau}$, we get the required result. □

Lemma 3.8. *Consider the coalescence process starting with set S_0, satisfying $|S_0| \leqslant 4\alpha\log\alpha$, where $\alpha = t_{\text{meet}}/t_{\text{mix}}$. Let $T_2 = \min\{t \geqslant 0 \mid |S_t| \leqslant 1\}$. Then $\mathbb{E}[T_2] = O(t_{\text{meet}})$.*

Proof. We will consider the coalescence process in phases. Let ℓ be the largest integer such that $|S_0| \geqslant \left(\frac{4}{3}\right)^\ell$. For $j \geqslant 1$, the j^{th} phase ends when $|S_t| < \left(\frac{4}{3}\right)^{\ell-j+1}$. The $(j+1)^{th}$ phase begins as soon as the j^{th} phase ends. Note that it may be the case that some phases are empty. Let $T_2(j)$ denote the time for phase j to last. We will only consider phases up to which $\ell - j + 1 \geqslant 32$.

Now we focus on a particular phase j. Let t_j be the time when the j^{th} phase begins and let S_{t_j} denote the corresponding set at that time. Thus, we have

$$\left(\frac{4}{3}\right)^{\ell-j+1} \leqslant |S_{t_j}| < \left(\frac{4}{3}\right)^{\ell-j+2}$$

We consider the process P^c defined in Section 2 as follows. If $n_j = |S_{t_j}|$ and $S_{t_j} = \{v_1, \ldots, v_{n_j}\}$, then $S_0^c = \{(v_1, 1), \ldots, (v_{n_j}, n_j)\}$. Then, according to the standard coupling $S_{t_j+t} = \{v \mid \exists i \in \mathbb{N}, (v, i) \in S_t^c\}$ with $t \in \mathbb{N}$. Thus, phase j ends when $|S_t^c| < \left(\frac{4}{3}\right)^{\ell-j+1}$. Let $S_0' = S_0^c$, let $k_j = \left\lceil \frac{|S_0'|}{2} \right\rceil$ and consider the process $P^*(S_0', k_j)$ as defined in Section 3.3. Let $g_j = \left\lfloor \frac{|S_0'|-k_j}{3} \right\rfloor$ and $T_2^*(j) = \min\{t \mid |\mathsf{IDs}(S_t') \cap G_2| < g_j\}$. We note that as long as $\ell - j + 1 \geqslant 32$, $g_j \geqslant 1$ and at time $T_2^*(j)$, $|S_t'| \leqslant \frac{3}{4} \cdot |S_0'|$. Since $\sigma^c \prec \sigma^*$ by Observation 3.2, $T_2^*(j)$ stochastically dominates $T_2(j)$ and hence it suffices to bound $\mathbb{E}[T_2^*(j)]$. In order to bound $\mathbb{E}[T_2^*(j)]$, we define

$$b_j = 32\alpha \log(4/3)(\ell - j + 1)(3/4)^{\ell-j+1}.$$

It can be verified that $1 \leqslant b_j \leqslant ((e-1)/e)\alpha$. Applying Lemma 3.4 with this value of b_j, we get that for $\tau_j = cb_j t_{\text{mix}}$, for independent random walks $(X_t)_{t=0}^{\tau_j}$, $(Y_t)_{t=0}^{\tau_j}$, intersect$(X_t, Y_t, \tau_j) \geqslant p_j$, where $p_j = 32\log(4/3)(\ell - j + 1)(3/4)^{\ell-j+1}$.

We will apply Claim 3.6 to bound $\mathbb{E}[T_2^*(j)]$. We need to verify that $k_j \geqslant \frac{8\log(4/p_j)}{p_j}$; to see this consider the following:

$$\frac{8\log(4/p_j)}{p_j} = \frac{8((\ell - j + 1)\log(4/3) - \log(8\log(4/3)(\ell - j + 1)))}{32\log(4/3)(\ell - j + 1)} \cdot \left(\frac{4}{3}\right)^{\ell-j+1}$$

For $\ell - j + 1 \geqslant 32$, it is the case that $(\ell - j + 1)\log(4/3) \geqslant \log(8\log(4/3)(\ell - j + 1))$. Thus, we have,

$$\frac{8\log(4/p_j)}{p_j} \leqslant \frac{16(\ell - j + 1)\log(4/3)}{32\log(4/3)(\ell - j + 1)} \cdot \left(\frac{4}{3}\right)^{\ell-j+1} \leqslant \frac{1}{2} \cdot |S_0'| \leqslant k_j \tag{6}$$

Since $b_j \geqslant 1$, there exists a suitably large constant c_1, so that $4t_{\text{mix}} + \tau_j \leqslant c_1 b_j t_{\text{mix}}$, using (6) we get that $\sqrt{\log(4/p_j)/p_j} \leqslant (\sqrt{4/3})^{\ell-j+1}$, and by definition $|G_2|/g \leqslant 3$, hence $\log|G_2| - \log g \leqslant \log(3)$, we get that there is a constant c_2 such that,

$$\mathbb{E}[T_2^*(j)] \leqslant c_2 \cdot t_{\text{mix}} \cdot \alpha \cdot (\ell - j + 1)\left(\frac{3}{4}\right)^{(\ell-j+1)/2} \tag{7}$$

Note that since we stop when $\ell - j + 1 < 32$, there are at most $\ell - 30$ phases considered. Let $\widetilde{T}$ be the random variable denoting the time-step when the last phase ends; at this point $|S_{\widetilde{T}}| = O(1)$. Therefore, using Proposition 3.3, $\mathbb{E}\left[T_2 - \widetilde{T} \,|\, \widetilde{T}\right] = O(t_{\text{meet}})$. But, clearly $\widetilde{T}$ is stochastically dominated by $\sum_{j=1}^{\ell-30} T_2^*(j)$. Thus, we have

$$\begin{aligned}
\mathbb{E}[T_2] &= \mathbb{E}\left[\widetilde{T}\right] + \mathbb{E}\left[\mathbb{E}\left[T_2 - \widetilde{T} \mid \widetilde{T}\right]\right] \\
&\leqslant c_2 \cdot t_{\text{mix}} \cdot \alpha \sum_{j=1}^{\ell-30} (\ell - j + 1)\left(\frac{3}{4}\right)^{(\ell-j+1)/2} + c_3 t_{\text{meet}} &(8)\\
&\leqslant c_2 \cdot t_{\text{mix}} \cdot \alpha + c_3 t_{\text{meet}} = O(t_{\text{meet}}) &(9)
\end{aligned}$$

Above, in (8) we used (7) and the fact that $\mathbb{E}\left[T_2 - \widetilde{T} \mid \widetilde{T}\right] \leqslant c_3 t_{\text{meet}}$ for some constant $c_3 > 0$. In step (9), we used the fact that $\sum_{j=1}^{\infty} jc^j = O(1)$ for $c < 1$. □

Thus, the first phase (Lemma 3.7) and the second phase (Lemma 3.8) take together $O(\sqrt{\alpha} \cdot \log n \cdot t_{\text{mix}} + t_{\text{meet}})$ time-steps, which yields Theorem 1.1.

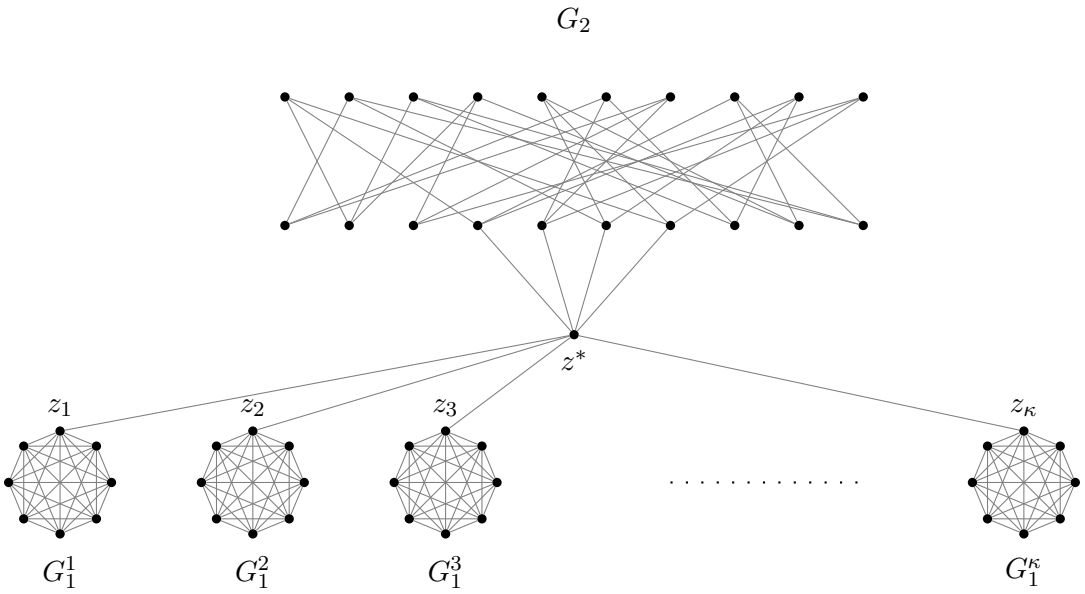

Figure 3: The graph described in Section 3.4 with $t_{\text{coal}} = \Omega(t_{\text{meet}} + \sqrt{\alpha} \cdot \log n \cdot t_{\text{mix}})$.

3.4 Lower Bound - Proof of Theorem 1.2

In this section we give a construction of a graph family in order to establish lower bounds on $t_{\text{coal}}(G)$ in terms of $t_{\text{meet}}(G)$ and $t_{\text{mix}}(G)$ demonstrating that Theorem 1.1 is asymptotically tight. Additionally, our construction generalizes a claim of Aldous and Fill [2, Chapter 14]: They mention that it is possible construct regular graphs that mimic the n-star in the sense that the $t_{\text{meet}} = o(t_{\text{avg-hit}})$, without giving further details of the construction. Our construction shows that even the coalescence time can be significantly smaller than the average hitting time for almost-regular graphs. For our family of almost-regular graphs, there is a polynomial gap between t_{meet} and $t_{\text{avg-hit}}$. More importantly, we show that these almost-regular graphs have a gap of $\sqrt{t_{\text{mix}}/t_{\text{meet}}} \log n$ between coalescing and meeting time. This shows that the bound in Theorem 1.1 is best possible, even if we constrain G to be almost-regular. We refer the reader to Section 1.2 for a high-level description of the proof ideas.

More precisely, in the proof of Theorem 1.2 we shall give an explicit construction of a graph family $G = G_n$ with $t_{\text{coal}} = \Omega(\sqrt{\alpha_n} \cdot \log n \cdot t_{\text{mix}})$, where $\alpha_n = \sqrt{t_{\text{meet}}/t_{\text{mix}}}$. For the remainder of this section, we will drop the dependence on n and will simply G instead of G_n and α instead α_n.

The construction of G (see Figure 3 for an illustration) is based on two building blocks, G_1 and G_2. First, let $G_1 = (V_1, E_1)$ be a clique of size $\sqrt{n}$. Let $G_2 = (V_2, E_2)$ be a $\sqrt{n}$-regular bipartite Ramanujan Graph on $n/\sqrt{\alpha'}$ nodes [26], where $\alpha' = \max\{\alpha, 240\}$. The graph G is made of one copy of G_2, $\kappa = \sqrt{n}$ copies of G_1 (denoted by $G_1^1, G_1^2, \ldots, G_1^\kappa$), and a node z^*, which has an edge to $\sqrt{n/\alpha'}$ distinct nodes of G_2 and to each of the designated nodes $z^i \in V_1^i$ in G_1^i for $i \in [1, \kappa]$. It is not difficult to see that this graph is almost-regular, *i. e.,* maximum and minimum degree differ by at most a constant.

In Lemma 3.13, Lemma 3.14, Lemma 3.15 and Lemma 3.16 respectively we show that $t_{\text{mix}} = \Theta(n)$, $t_{\text{meet}} = \Theta(\alpha' n)$, $t_{\text{coal}} = \Omega(\sqrt{\alpha'} \cdot n \log n)$, and $t_{\text{avg-hit}} = \Omega(n^{3/2})$. We start with the following auxiliary lemma which shows that the walk restricted to V_2 behaves similarly to the walk restricted to $V_2 \cup \{z^*\}$, meaning that the walks have very similar t-step probabilities.

Lemma 3.9. *Let Q be the transition matrix of the random walk on G_2 and $\widehat{Q}$ be the transition matrix of the random walk on the subgraph induced by $V_2 \cup \{z^*\}$. Let $S^* = \{u \in V_2 \cap N(z^*)\}$. Then the following statements hold:*

(i) For any $u, v \in V_2$ we have $\|p^t_{u,\cdot} - q^t_{u,\cdot}\|_{TV} \leqslant \sum_{i=1}^{t-1} p^i_{u,S^*}/(2\sqrt{n}) \leqslant t/(2\sqrt{n})$.

(ii) For any $u, v \in V_2$ we have $\|\widehat{q}^t_{u,\cdot} - q^t_{u,\cdot}\|_{TV} \leqslant \sum_{i=1}^{t-1} p^i_{u,S^*}/(2\sqrt{n}) \leqslant t/(2\sqrt{n})$.

(iii) For any $u, v \in V_2$ we have that after $t = t_{\mathrm{mix}}(G_2)$ time steps $\|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{TV} \leqslant o(1) + 2/e$.

Proof. Let $(X_t)_{t\geqslant 0}$ be the Markov chain with transition matrix P and let $(Y_t)_{t\geqslant 0}$ be the Markov chain with transition matrix Q. We will inductively couple these two random walks starting from $X_0 = Y_0 = u$. Given that we coupled both chains up to time $t-1$, we can couple (X_t, Y_t) such that the coupling fails w.p.

$$\begin{aligned}\mathbb{P}[X_t \neq Y_t \mid X_{t-1} = Y_{t-1}] &= \mathbb{P}[X_t \neq Y_t \mid X_{t-1} = Y_{t-1}, X_{t-1} \in S^*] \cdot \mathbb{P}[X_{t-1} \in S^*] \\ &\quad + \mathbb{P}[X_t \neq Y_t \mid X_{t-1} = Y_{t-1}, X_{t-1} \in V_2 \setminus S^*] \cdot \mathbb{P}[X_{t-1} \in V_2 \setminus S^*] \\ &\leqslant p^{t-1}_{u,S^*}/(2\sqrt{n}) + 0.\end{aligned}$$

We have, by [23, Proposition 4.7], $\|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\mathrm{TV}} = \inf\{\mathbb{P}[X \neq Y] \mid (X, Y) \text{ is a coupling of } p^t_{u,\cdot} \text{ and } p^t_{v,\cdot}\}$. Hence, by a union bound over t steps,

$$\begin{aligned}\|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\mathrm{TV}} &= \inf\{\mathbb{P}[X \neq Y] \mid (X, Y) \text{ is a coupling of } p^t_{u,\cdot} \text{ and } p^t_{v,\cdot}\} \leqslant \mathbb{P}[X_t \neq Y_t] \\ &\leqslant \sum_{i=1}^{t-1} p^i_{u,S^*}/(2\sqrt{n}) \leqslant \frac{t}{2\sqrt{n}}.\end{aligned}$$

To prove the second part we redefine $(X_t)_{t\geqslant 0}$ to be the Markov chain with transition matrix $\widehat{Q}$ and the proof follows the same lines.

We proceed with the last part. For $u, v \in V_2$ we have that after $t = t_{\mathrm{mix}}(G_2)$ time steps, by the triangle inequality, $\|p^t_{u,\cdot} - \pi^{G_2}(\cdot)\|_{\mathrm{TV}} \leqslant \|p^t_{u,\cdot} - q^t_{u,\cdot}\|_{\mathrm{TV}} + \|q^t_{u,\cdot} - \pi^{G_2}(\cdot)\|_{\mathrm{TV}} \leqslant o(1) + 1/e$. Again, by the triangle inequality, $\|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\mathrm{TV}} \leqslant o(1) + 2/e$. □

Based on Lemma 3.9, we can now bound the hitting time to reach z^*, which will later be used to establish the bounds on the mixing and meeting time of the whole graph G. But first, we prove that the mixing time of the graph G' induced by $V_2 \cup \{z^*\}$ is constant and that after mixing on G', the random walk has a probability of $\Omega(1/n)$ to hit z^* in two time-steps.

Lemma 3.10. *The following three statements hold.*

(i) Let G' be the induced graph by the vertices $V_2 \cup \{z^*\}$. Then $t_{\mathrm{mix}}(G') = O(1)$.

(ii) Let $u \in V \setminus \{z^*\}$. Then there exists a constant $c \geqslant 1$ such that $\mathbb{P}[\tau_{u,z^*} \geqslant n/c] \geqslant 1/2$.

(iii) Let $u \in V \setminus \{z^*\}$. Then $t_{\mathrm{hit}}(u, z^*) = O(n)$.

Proof. We prove the statements one by one.

(i) Let Q be the transition matrix of a random walk restricted to G_2. Let $d^Q(t)$ be the total variation distance w.r.t. the transition matrix Q. Further, let $\widehat{Q}$ be the transition matrix of a random walk restricted to G'. Observe that $t_{\mathrm{mix}}(G_2) = O(1)$, by Lemma A.2.

Let $t \in [2t_{\text{mix}}(G_2), \log n]$. Let $\mathcal{D}$ be the set of distributions such that for every $D' \in \mathcal{D}$ and $u \sim D'$ we have $\mathbb{P}[\, u = z^* \,] = 0$. We have

$$\|q^t_{u\sim D',\cdot} - \widehat{\pi}^{G'}(\cdot)\|_{\text{TV}} = \frac{1}{2}\sum_{v\in V_2}\left|\widehat{q}^t_{u\sim D',v} - \pi^{G'}(v)\right| + \frac{1}{2}\left|\widehat{q}^t_{u\sim D',z^*} - \pi^{G'}(z^*)\right|$$

For $u \in V_2$ observe that $\pi^{G'}(u) \in [\pi^{G_2}(u)(1-\varepsilon), \pi^{G_2}(u)(1+\varepsilon)]$ for some $\varepsilon = o(1)$. Thus, for $\delta_v := |\widehat{q}^t_{u,v} - q^t_{u,v}|$, we get by using triangle inequality,

$$\begin{aligned}
\frac{1}{2}\sum_{v\in V_2}\left|\widehat{q}^t_{u\sim D',v} - \pi^{G'}(v)\right| &\leqslant \frac{1}{2}\sum_{v\in V_2}\left|\widehat{q}^t_{u\sim D',v} - \pi^{G_2}(v)\right| + \frac{1}{2}\sum_{v\in V_2}\pi^{G_2}(v)|\varepsilon| \\
&\leqslant \frac{1}{2}\sum_{v\in V_2}\left|q^t_{u\sim D',v} - \pi^{G_2}(v)\right| + \frac{1}{2}\sum_{v\in V_2}|\delta_v| + \frac{1}{2}\sum_{v\in V_2}\pi^{G_2}(v)|\varepsilon| \\
&\leqslant d^Q(t) + 1/32 + 1/32,
\end{aligned}$$

since, by Lemma 3.9, $\frac{1}{2}\sum_{v\in V}|\delta_v| \leqslant t/(2\sqrt{n}) \leqslant \frac{1}{32}$. Putting everything together and using $\frac{1}{2}|\widehat{q}^t_{u\sim D',z^*} - \pi^{G'}(z^*)| \leqslant \log n/\sqrt{n} \leqslant 1/32$, we get

$$\begin{aligned}
\|\widehat{q}^t_{u\sim D',\cdot} - \pi^{G'}(\cdot)\|_{\text{TV}} &= \frac{1}{2}\sum_{v\in V_2}\left|\widehat{q}^t_{u\sim D',v} - \pi^{G'}(v)\right| + \frac{1}{2}\left|\widehat{q}^t_{u\sim D',z^*} - \pi^{G'}(z^*)\right| \\
&\leqslant d^Q(t) + 1/32 + 1/32 + 1/32 \leqslant 1/e^2 + 3/32 \leqslant 1/e,
\end{aligned}$$

where the second last inequality follows by the definition of the mixing time together with [23, Exercise 4.1] giving that $d(t) = \max_{D\in\mathcal{D}^*}\|p^t_{u\sim D,\cdot} - \pi(\cdot)\|_{\text{TV}}$, where $\mathcal{D}^*$ ranges over all distributions over $V(G')$.

Furthermore, observe that $\widehat{q}^t_{z^*,z^*} \leqslant 1/2^t + \sum_{i\leqslant t}\sum_{v\in N(z^*)} q^{i-1}_{z^*,v}\cdot 1/\sqrt{n} \leqslant 1/2^t + t/\sqrt{n} \leqslant 1/32$ for $t = 7$.

Consider the random walk starting at z^* and let $(X_0, X_1, \dots)$ denote its trajectory. Let $D' \in \mathcal{D}$ denote the probability distribution of the position of the random walk at time-step 7 conditioned on $X_7 \neq z^*$, *i. e.,* D' is the probability distribution given by $\mathbb{P}[\, X_7 = u \mid u \in V_2 \,]$. Hence,

$$\|\widehat{q}^{t+7}_{z^*,\cdot} - \pi^{G'}(\cdot)\|_{\text{TV}} \leqslant q_{z^*,V(G')\setminus\{z^*\}}\cdot\|\widehat{q}^t_{u\sim D',\cdot} - \pi^{G'}(\cdot)\|_{\text{TV}} + q_{z^*,z^*}\cdot 1 \leqslant (1/e^2 + 3/32) + 1/32 \leqslant 1/e.$$

We conclude that for all $u \in V'$, $\|\widehat{q}^t_{u,\cdot} - \pi^{G'}(\cdot)\|_{\text{TV}} \leqslant 1/e$ and thus $t_{\text{mix}}(G') = O(1)$.

(ii) To prove $\mathbb{P}[\, \tau_{u,z^*} \geqslant n/c \,]$ for $u \in V_2$ we show that the random walk restricted to G' doesn't hit z^* after n/c_1 steps w.c.p. for some large enough constant c_1. By Union bound, for some large constants c_1, c_2 that

$$\begin{aligned}
\mathbb{P}\left[\, \tau^G(z^*) \leqslant n/c_1 \,\right] = \mathbb{P}\left[\, \tau^{G'}(z^*) \leqslant n/c_1 \,\right] &\leqslant \sum_{t=1}^{n/c_1}\widehat{q}_t(u, z^*) \leqslant \sum_{t=1}^{c_2\log n} 1/\sqrt{n} + \sum_{t=c_2\log n}^{n/c_1}\widehat{q}_t(u, z^*) \\
&\leqslant o(1) + n/c_1\cdot(\pi^{G'}(z^*) + 1/n^2) \leqslant 1/2,
\end{aligned}$$

where we used $\widehat{q}^t_{u,z^*} \leqslant \pi^{G'}(z^*) + \sqrt{\frac{\pi^{G'}(z^*)}{\pi^{G'}(u)}}\lambda_2(G')^t$ (Lemma A.1).

We proceed by showing that $\mathbb{P}\left[\tau_{u,z^*} \geqslant n/c\right]$ for $u \in V_1$. Consider first a random walk $(\widetilde{X}_t)_{t\geqslant 0}$ restricted to $G_1^1 = G_1$ that starts at vertex z. Furthermore, in order to couple the random walk $\widetilde{X}_t$ restricted to G_1 with a random walk in G, we will consider the random variable $\widetilde{Z} := \sum_{t=0}^{t_{\text{sep}}^{G_1}} \mathbf{1}_{\widetilde{X}_t = z_1}$. Since G_1 is a clique, $t_{\text{sep}}^{G_1} = O(1)$, and $p_{z_1,z_1}^t \leqslant \frac{1}{\sqrt{n}} + \lambda_2(G_1)^t$ by Lemma A.1, where $\lambda_2(G_1)$ is some constant bounded away form 1. Therefore, $\mathbb{E}\left[\widetilde{Z}\right] = \sum_{t=0}^{n/c_1} p_{z_1,z_1}^t \leqslant 2\sqrt{n}/c_1$. Let $\gamma := 4 \cdot \mathbb{E}\left[\widetilde{Z}\right]$. Then, by Markov's inequality

$$\mathbb{P}\left[\widetilde{Z} \geqslant \gamma\right] \leqslant 1/4.$$

Consider now the straightforward coupling between a random walk $(X_t)_{t\geqslant 1}$ in G that starts at vertex z^1 and the random walk $(\widetilde{X}_t)_{t\geqslant 1}$ restricted to G_1^i that starts at the same vertex. Whenever the random walk $\widetilde{X}_t$ is at a vertex different from z^1, then the random walk X_t makes the same transition. If the random walk $\widetilde{X}_t$ is at vertex z^1, then there is a coupling so that the random walk X_t makes the same transition as $\widetilde{X}_t$ with probability $\frac{2\sqrt{n}-1}{2\sqrt{n}}$. Conditional on the event $\widetilde{Z} \leqslant \gamma$ occurring, the random walk $\widetilde{X}_t$ follows the random walk X_t up until step n/c_1 with probability at least

$$p := \left(\frac{2\sqrt{n}-1}{2\sqrt{n}}\right)^{\gamma} \geqslant 3/4,$$

since the random walk $\widetilde{X}_t$ has at most γ visits to z^1. Therefore, by Union bound,

$$\mathbb{P}\left[\tau^G(z^*) \geqslant n/c_1\right] \geqslant \mathbb{P}\left[X_t = \widetilde{X}_t, t \leqslant n/c_1\right] \geqslant 1 - \mathbb{P}\left[\widetilde{Z} \geqslant c_3\right] - (1-p) \geqslant 1/2$$

and the proof is complete.

(iii) We proceed by showing $t_{\text{hit}}(u, z^*) = O(n)$ for $u \in V_2$.

Let Q be the transition matrix of the random walk restricted to G_2. Let $u \in V_2$ and $S^* = N(z^*)$ be the neighbors of z^* in G_2. Furthermore, After $t = t_{\text{sep}}(G_2)$ we have that

$$q_{u,S^*}^t := \sum_{v\in S^*} q_{u,v}^t \geqslant \sum_{v\in S^*} \pi^{G_2}(v)(1-e^{-1}) \geqslant \frac{\sqrt{n}}{\sqrt{\alpha'}} \cdot \frac{\sqrt{\alpha'}}{1.2n}(1-e^{-1}) = \frac{1-e^{-1}}{1.2\sqrt{n}}.$$

By Lemma 3.9, we have for any $u \in V_2$ that $\|p_{u,\cdot}^t - q_{u,\cdot}^t\|_{\text{TV}} \leqslant t_{\text{sep}}(G_2)/(2\sqrt{n})$. To bound $\tau^G(z^*)$ we show that after $O(1)$ steps the random walk hits z^* w.p. $\Omega(1/n)$.

We distinguish between two cases.

(a) For all $i \leqslant t$ we have $p_{u,S^*}^t \leqslant 1/t_{\text{sep}}(G_2)$. Thus,

$$p_{v,S^*}^t = \sum_{v\in S^*} p_{u,v}^t \geqslant \frac{1-e^{-1}}{1.2\sqrt{n}} - \frac{t_{\text{sep}}(G_2)}{t_{\text{sep}}(G_2)2\sqrt{n}} = \Omega(1/\sqrt{n}).$$

Hence, the random walk hits z^* after $t_{\text{sep}}(G_2)+1$ w.p. at least $p_1 := p_{u,S^*}^t \cdot p_{u\in S^*,z^*} = \Omega(1/n)$.

(b) Otherwise there exists a t^* such that $p_{u,S^*}^{t^*} > 1/t_{\text{sep}}(G_2)$. Thus the random walk hits z^* after $t_{\text{sep}}(G_2)+1$ w.p. at least $p_2 := \min_{u\in S^*}\{p_{u,S^*}^{t^*} \cdot p_{u,z^*}\} = \Omega(1/\sqrt{n})$.

Thus after $O(1)$ steps the random walk hits z^* w.p. $\Omega(1/n)$.

We now show a similar statement if $u \in V_1$. Let $(X_t)_{t\geqslant 0}$ be the random walk starting on u. Observe that X_t hits z^* with probability $p_1(u,z^1)\cdot p_1(z^1,z^*) = \Omega(1/n)$ in 2 time steps. Hence, for any $u \in V$ we $\mathbb{P}[\,\tau_u(z^*) = O(1)\,] = \Omega(1/n)$. Thus, repeating this iteratively and using independence yields $t_{\text{hit}}(u,z^*) = O(n)$ for $u \in V$.

□

To establish the mixing time of G, we will make use of the following result of Peres and Sousi.

Theorem 3.11 ([31]). *For any $\beta < 1/2$, let $t_H(\beta) = \max_{u,A:\pi(A)\geqslant\beta} t_{\text{hit}}(u,A)$. Then there exist positive constants c_β and c'_β such that*

$$c'_\beta \cdot t_H(\beta) \leqslant t_{\text{mix}}(1/4) \leqslant c_\beta \cdot t_H(\beta).$$

In the following we show for any β close enough to $1/2$, that any $A \subseteq V$ satisfying $\pi(A) \geqslant \beta$ must include at least a constant fraction of nodes from a constant fraction of copies of G_1.

Claim 3.12. *Let $\beta = 1/2 - 10^{-3}$. For any $A \subseteq V$ with $\pi(A) \geqslant \beta$, define $H(A) = \{i \mid |G_1^i \cap A| \geqslant |V_1|/(2e)\}$. Then, $|H(A)| \geqslant \kappa/(2e)$.*

Proof. This follows from a simple pigeon hole argument: Suppose $|H(A)| < \kappa/(2e)$ was true. Then,

$$\begin{aligned}\pi(A) &\leqslant |H(A)|\cdot\pi(V_1) + (\kappa - |H(A)|)\cdot\left(\frac{\pi(V_1)}{2e} + \pi(z^i)\right) + \pi(V_2) + \pi(z^*)\\ &< \frac{\kappa}{2e}\cdot\pi(V_1) + \kappa\cdot\left(\frac{\pi(V_1)}{2e} + \pi(z^i)\right) + 1/20 < \beta \leqslant \pi(A),\end{aligned}$$

which is a contradiction and hence choice of A must fulfill $|H(A)| \geqslant \kappa/(2e)$. □

We are now ready to determine the mixing time of G. The lower bound is a simple application of Cheeger's inequality, while the upper bound combines the previous lemmas with Theorem 3.11.

Lemma 3.13. *Let G be the described at the beginning of Section 3.4. We have $t_{\text{mix}}(G) = \Theta(n)$.*

Proof. First we show $t_{\text{mix}} = \Omega(n)$. The *conductance* of $G = (V,E)$ is defined by $\Phi(G) = \min_{\substack{U\subseteq V,\\ 0<\text{vol}(U)\leqslant\text{vol}(V)/2}} \frac{|E(U,V\setminus U)|}{\text{vol}(U)}$. In particular, for $U = V_1$ we get that $\Phi(G) \leqslant \frac{4}{n}$. Hence, by Cheeger's inequality and $\left(\frac{1}{1-\lambda_2(G)} - 1\right)\cdot\log(\frac{e}{2}) \leqslant t_{\text{mix}}(1/e)$ (see, *e. g.*, [23, Chapter 12]),

$$\frac{n}{4} \leqslant \frac{1}{\Phi(G)} \leqslant \frac{2}{1-\lambda_2(G)} = \frac{2}{1-\lambda_2(G)} - 2 + 2 \leqslant \frac{2t_{\text{mix}}}{\log\left(\frac{e}{2}\right)} + 2.$$

Rearranging the terms yields $t_{\text{mix}} = \Omega(n)$.

We proceed with the upper bound on the mixing time. Let $\beta = 1/2 - 10^{-3}$ and let $A \subseteq V$ be an arbitrary set satisfying $\pi(A) \geqslant \beta$. First, we apply Claim 3.12 to conclude that $|H(A)| \geqslant \kappa/(2e)$. This immediately implies that with $Z := \{z^i \colon i \in H(A)\}$, $|Z| \geqslant \kappa/(2e)$. The remainder of the proof is divided into the following three parts:

(i) Starting from any vertex $u \in V$, with probability at least $1/2$, the random walk hits z^* after $2\max_{u\in V} t_{\text{hit}}(u,z^*) = O(n)$ steps.

(ii) With constant probability $p_1 > 0$, the random walk moves from z^* to a vertex in Z.

(iii) With constant probability $p_2 > 0$ a random walk starting from a vertex in Z will hit A after one step.

It is clear that combining these three results shows that with constant probability $\frac{1}{2}p_1p_2 > 0$, a random walk starting from an arbitrary vertex $u \in V$ hits a vertex in A after $O(n)+1+1$ time-steps. Iterating this and using independence shows that $t_{\text{hit}}(u, A) = O(n)$, and hence by Theorem 3.11, $t_{\text{mix}} = O(n)$ as needed.

Part (i). Consider $\max_{u\in V} t_{\text{hit}}(u, z^*)$. For $u \in V$, Lemma 3.10.(*iii*) implies $t_{\text{hit}}(u, z^*) = O(n)$.

Part (ii). If the random walk is on z^*, then since $\deg(z^*) = \kappa + \sqrt{n/\alpha'}$, $|Z| \geqslant \kappa/(2e)$, it follows that the random walk hits a vertex in Z after one step with constant probability $p_1 := \frac{|Z|}{2(\kappa+\sqrt{n/\alpha'})} > 0$.

Part (iii). Finally, for any $z \in Z$ we have that $p_2 = p(z, A) = \frac{|V_1|/(2e)}{2\sqrt{n}} > 0$ and the proof is complete. □

In the following we establish the bound on the meeting time. As it turns out, any meeting is very likely to happen on V_2 and it takes about $\Theta(\alpha' n)$ time-steps until both walks reach V_2 simultaneously. The lower bound then follows from our common analysis method (1). The upper bound combines the mixing time bound of $O(n)$ (Lemma 3.13), and that once a random walk reaches V_1, it says there for $\Theta(n)$ steps with constant probability Lemma 3.10.ii.

Lemma 3.14. *Let G be the described at the beginning of Section 3.4. We have $t_{\text{meet}}(G) = \Theta(\alpha' n)$.*

Proof. We start by proving $t_{\text{meet}} = \Omega(\alpha' n)$: Consider two non-interacting, random walks with starting positions drawn from the stationary distribution π. Let $\ell = c'\alpha' n$, for some small enough constant c'. Let Z_1 be the number of collisions of the two random walks on the nodes in $V_1^1 \cup V_1^2 \cup \cdots \cup V_1^\kappa$. Let Z_2 be the number of collisions of the two random walks on the nodes in V_2. Let Z_* be the number of collisions of the two random walks on the node z^*.

Let Z be the number of collisions of the two walks during the first ℓ time steps, *i. e.*, $Z = Z_1 + Z_2 + Z_*$. Using union bound we derive

$$\begin{aligned}\mathbb{P}[\,Z \geqslant 1\,] &\leqslant \mathbb{P}[\,Z_1 \geqslant 1\,] + \mathbb{P}[\,Z_2 \geqslant 1\,] + \mathbb{P}[\,Z_* \geqslant 1\,] \\ &\leqslant \frac{\mathbb{E}[\,Z_1\,]}{\mathbb{E}[\,Z_1 \,|\, Z_1 \geqslant 1\,]} + \frac{\mathbb{E}[\,Z_2\,]}{\mathbb{E}[\,Z_2 \,|\, Z_2 \geqslant 1\,]} + \frac{\mathbb{E}[\,Z_*\,]}{\mathbb{E}[\,Z_* \,|\, Z_* \geqslant 1\,]}. \qquad (10)\end{aligned}$$

We have $\mathbb{E}[\,Z_1\,] \leqslant \ell n \left(\frac{2}{n}\right)^2$, $\mathbb{E}[\,Z_2\,] \leqslant \ell \frac{n}{\sqrt{\alpha'}} \left(\frac{2}{n}\right)^2$, and $\mathbb{E}[\,Z_*\,] \leqslant \ell \left(\frac{2}{n}\right)^2$, since $\max_u \pi(u) \leqslant 2/n$. Conditioning on $Z_1 \geqslant 1$ and since both random walks start from the stationary distribution, we have, by Observation A.6, that the first meeting happens in the first $\ell/2$ time steps w.p. at least $1/2$.

Consider $\mathbb{E}[\,Z_1 \,|\, Z_1 \geqslant 1\,]$. Suppose the meeting occurred at node $u \in V_1$. Let $\mathcal{E}_1$ be the event that $\tau_{u,z^*} \geqslant n/c$ for both walks, where $c > 0$ is a large enough constant. By Lemma 3.10.(ii), we have that $\mathbb{P}[\,\mathcal{E}_1\,] \geqslant 1/2$. Let $\widehat{p}^t_{u,\cdot}$ for $t \leqslant n/c$ be the distribution of the random walk starting on u after t time steps under the conditioning $\mathcal{E}_1$. Observe that $\sum_{v\in V_1} \widehat{p}^t_{u,v} = 1$ implying that $\sum_{v\in V_1} (\widehat{p}^t_{u,v})^2 \geqslant \sum_{v\in V_1} \left(\frac{1}{|V_1|}\right)^2 = 1/|V_1|$. Hence, since $t/2 \geqslant n/c$ we get

$$\mathbb{E}[\,Z_1 \,|\, Z_1 \geqslant 1\,] \geqslant \mathbb{E}[\,Z_1 \,|\, Z_1 \geqslant 1, \mathcal{E}_1\,] \cdot \mathbb{P}[\,\mathcal{E}_1\,] \geqslant \frac{1}{2} \min_{u\in V_1} \sum_{t=0}^{n/c} \sum_{v\in V_1} (\widehat{p}^t_{u,v})^2 \geqslant \frac{1}{2} \sum_{t=0}^{n/c} 1/|V_1| = \frac{\sqrt{n}}{2c}.$$

Using an exactly analogous analysis for Z_2 we can upper bound $\mathbb{E}[Z_2 \mid Z_2 \geqslant 1]$ as follows:

$$\mathbb{E}[Z_2 \mid Z_2 \geqslant 1] \geqslant \mathbb{E}[Z_2 \mid Z_2 \geqslant 1, \mathcal{E}_2] \cdot \mathbb{P}[\mathcal{E}_2] \geqslant \frac{1}{2} \min_{u \in V_2} \sum_{t=0}^{n/c} \sum_{v \in V_2} (\widehat{p}^t_{u,v})^2 \geqslant \frac{1}{2} \sum_{t=0}^{n/c} 1/|V_2| = \frac{\sqrt{\alpha'}}{2c},$$

where $\mathcal{E}_2$ is the event that $\tau_{u,z^*} \geqslant n/c$ for some large enough constant c and $u \in V_2$. Plugging everything into (10) yields

$$\begin{aligned}
\mathbb{P}[Z \geqslant 1] &\leqslant \frac{\mathbb{E}[Z_1]}{\mathbb{E}[Z_1 \mid Z_1 \geqslant 1]} + \frac{\mathbb{E}[Z_2]}{\mathbb{E}[Z_2 \mid Z_2 \geqslant 1]} + \frac{\mathbb{E}[Z_*]}{\mathbb{E}[Z_* \mid Z_* \geqslant 1]} \\
&\leqslant \frac{\ell n \left(\frac{2}{n}\right)^2}{\frac{\sqrt{n}}{2c}} + \frac{\ell \frac{n}{\sqrt{\alpha'}} \left(\frac{2}{n}\right)^2}{\frac{\sqrt{\alpha'}}{2c}} + \frac{\ell \left(\frac{2}{n}\right)^2}{1} \\
&\leqslant o(1) + 8c \cdot c' + o(1) \leqslant 1/2,
\end{aligned}$$

for any constant $c' \in (0, \frac{1}{17c}]$. This finishes the proof of $t_{\text{meet}} = \Omega(\alpha' n)$.

In the remainder we prove $t_{\text{meet}} = O(\alpha' n)$. Consider two independent walks $(X_t)_t$ and $(Y_t)_t$, both starting from arbitrary nodes. Note $t_{\text{sep}} = O(t_{\text{mix}}) = O(n)$ by Lemma 3.13, and

$$\mathbb{P}\left[\left\{X_{t_{\text{sep}}} \in V_2\right\} \cap \left\{Y_{t_{\text{sep}}} \in V_2\right\}\right] \geqslant \left(\sum_{u \in V_2} (1-e)\pi(u)\right)^2 = \Omega\left(\left(1/\sqrt{\alpha'}\right)^2\right) = \Omega\left(1/\alpha'\right).$$

From Lemma 3.9 we get for $t = t_{\text{mix}}(G_2)$ that $\|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\text{TV}} \leqslant o(1) + 2/e$.

Hence $\sum_{w \in V} \min\{p^{t_{\text{mix}}(G_2)}_{u,w}, p^{t_{\text{mix}}(G_2)}_{v,w}\} \geqslant 1 - (e/2 + o(1)) \geqslant 1/10$ and thus $\sum_{w \in V} p^{t_{\text{mix}}(G_2)}_{u,w} p^{t_{\text{mix}}(G_2)}_{v,w} \geqslant 1/(100n)$.

We next divide the random walk into n consecutive trials of length $t_{\text{mix}}(G_2)$ each. For each trial, we have a random variable $X_i \in \{-1, 0, +1\}$, where $X_i = 1$ if the two random walks are on the same vertex in V at the end of the trial, $X_i = -1$ if at least one of the two random walk is on a vertex outside G_2 at the end of the trial and $X_i = 0$ otherwise. From the above, we have $\mathbb{P}\left[X_i = 1 \mid \cup_{j=1}^{i-1} \{X_i = 0\}\right] \geqslant c_1/n$ and $\mathbb{P}\left[X_i = -1 \mid \cup_{j=1}^{i-1} \{X_i = 0\}\right] \leqslant c_2/n$ for constants $c_1 > 0$ and $c_2 > 0$.

Let $\tau := \min\{i \colon X_i \neq 0\}$ be the first trial when the two random walks meet or one of the walks leaves G_2. Then by definition of the process, $\mathbb{P}[\tau \leqslant n] > c_3$, where $c_3 > 0$ is another constant. Furthermore for any integer σ, $\mathbb{P}[X_\tau = 1 \mid \tau = \sigma] \geqslant \frac{c_1/n}{c_1/n + c_2/n} = \frac{c_1}{c_1 + c_2}$. Therefore the probability that the two random walks have met after $O(n)$ trials is at least

$$\mathbb{P}[\{\tau \leqslant n\} \cap \{X_\tau = 1\}] \geqslant \mathbb{P}[\tau \leqslant n] \cdot \mathbb{P}[X_\tau = 1 \mid \tau \leqslant n] \geqslant c_3 \cdot \frac{c_1}{c_1 + c_2}.$$

Putting everything together, we get that two random walks meet in time $t_{\text{sep}}(G) + n \cdot t_{\text{mix}}(G_2) = O(n)$ w.p. at least $\mathbb{P}\left[\left\{X_{t_{\text{sep}}} \in V_2\right\} \cap \left\{Y_{t_{\text{sep}}} \in V_2\right\}\right] \cdot \mathbb{P}[\{\tau \leqslant n\} \cap \{X_\tau = 1\}] = \Omega(1/\alpha')$. Repeating this iteratively and using independence yields $t_{\text{meet}} = O(\alpha' n)$. ☐

Finally, we analyze the coalescing time of G. The proof idea is to consider $\sqrt[5]{n}$ random walks starting from π and show that meetings only occur on V_2 and that at least one random walk requires $\Omega(\sqrt{\alpha'} \cdot n \log n)$ time-steps to reach V_2.

Lemma 3.15. *Let G be the described at the beginning of Section 3.4. We have $t_{\text{coal}}(G) = \Omega(\sqrt{\alpha'} \cdot n \log n)$.*

Proof. Let $\varepsilon = 1/5$. We use the stochastic domination given by Observation 3.2 and show that even the coalescing time of n^ε random walks requires $\Omega(\sqrt{\alpha'} \cdot n \log n)$ time-steps w.c.p.. Let R be a collection of n^ε independent, *i. e.,* non-interacting, random walks with starting positions drawn from the stationary distribution π. We define the following three bad events:

(i) Let $\mathcal{E}_1$ be the event that any of the n^ε random walks meet on a node $V \setminus V_2$ in $\sqrt{\alpha'} \cdot n(\log n)^2$ steps.

(ii) Let $\mathcal{E}_2$ be the event that fewer than $n^\varepsilon/4$ random walks start on copies of G_1, *i. e.,* on nodes $V \setminus (V_2 \cup z^*)$.

(iii) Let $\mathcal{E}_3$ be the event that all random walks starting from V_1 require fewer than $c \cdot \sqrt{\alpha'} \cdot n \log n$ time-steps for leaving $V \setminus (V_2 \cup z^*)$ for some constant $c > 0$ to be determined later.

In the following we show that $\mathbb{P}[\mathcal{E}_1] = o(1)$, $\mathbb{P}[\mathcal{E}_2] = o(1)$, and $\mathbb{P}\left[\mathcal{E}_3 \mid \overline{\mathcal{E}_2}\right] < 1/e$, which implies, by union bound,

$$\mathbb{P}\left[\overline{\mathcal{E}_1} \cap \overline{\mathcal{E}_2} \cap \overline{\mathcal{E}_3}\right] \geqslant \mathbb{P}\left[\overline{\mathcal{E}_1}\right] - (1 - \mathbb{P}\left[\overline{\mathcal{E}_2} \cap \overline{\mathcal{E}_3}\right]) \geqslant 1 - o(1) - \left(1 - (1 - o(1)) \cdot \left(1 - \frac{1}{e}\right)\right) \geqslant 1 - \frac{1}{2e}.$$

Conditioning on $\overline{\mathcal{E}_1} \cap \overline{\mathcal{E}_2} \cap \overline{\mathcal{E}_3}$, none of the independent random walks meet on any node $V \setminus V_2$ and hence they are indistinguishable from coalescing random walks until they reach V_2. Therefore, it is necessary for all random walks to reach G_2 in order to coalesce. Hence, we conclude that $t_{\text{coal}}(G) = \Omega(\sqrt{\alpha'} \cdot n \log n)$ yielding the lemma.

(i) We now prove $\mathbb{P}[\mathcal{E}_1] = o(1)$. Consider any pair of the random walks R. Since both random walks start from the stationary distribution, the probability for them to meet on a node on z^* in a fixed step $t \geqslant 0$ is at most $O(1/n^2)$.

Hence, by the Union bound over $\binom{n^\varepsilon}{2}$ pairs of random walks and $\sqrt{\alpha'} \cdot n(\log n)^2$ steps, the probability of any two random walks meeting on z^* is at most

$$p_1 := \binom{n^\varepsilon}{2} \cdot \sqrt{\alpha'} \cdot n(\log n)^2 \cdot O(1/n^2) = o(1),$$

since $\varepsilon = \frac{1}{5}$. Furthermore, the probability that no two walks start on the same copy of G_1 is at most $p_2 := n^\varepsilon \cdot \frac{n^\varepsilon}{\sqrt{n}} = o(1)$, where we used union bound.

Moreover, using a Chernoff bound together with Lemma 3.10.(ii), it follows that a random walks visits the vertex z^* at most $10 \log n$ times during $n \log n$ steps with probability at least $1 - n^{-2}$. By the Union bound over all random walks, it follows that w.p. at least $1 - n^{-1}$, each random walk visits at most $10 \log n$ copies of G_1, and each such copy is chosen uniformly and independently at random among $G_1^1, G_1^2, \ldots, G_1^\kappa$. Therefore, the probability that there exists a copy of G_1 which is visited by at least two random walks in $n \log n$ steps is at most

$$p_3 := n^{-1} + n^\varepsilon(10 \log n + 1) \cdot \frac{n^\varepsilon(10 \log n + 1)}{\sqrt{n}} = o(1). \tag{11}$$

Putting everything together, using union bound, yields $\mathbb{P}[\mathcal{E}_1] \leqslant p_1 + p_2 + p_3 = o(1)$.

(ii) We now prove $\mathbb{P}[\mathcal{E}_2] = o(1)$. The probability p for each random walk to start on a node of $V \setminus (V_2 \cup z^*)$ is $\pi(V \setminus (V_2 \cup z^*)) \geqslant 1/2$. For each of the random walks i we define the indicator variable X_i to be one, if that random walk started on $V \setminus (V_2 \cup z^*)$. Let $X = \sum_{i=1}^{n^\varepsilon} X_i$. We

have $\mathbb{E}[X] = n^{\varepsilon} \cdot \mathbb{E}[X_i] \geqslant n^{\varepsilon}/2$. Since the starting positions of the n^{ε} random walks are drawn independently, by a Chernoff bound

$$\mathbb{P}[\mathcal{E}_2] = \mathbb{P}\left[X \leqslant \frac{1}{4}n^{\varepsilon}\right] \leqslant \mathbb{P}[X \leqslant \mathbb{E}[X]/2] \leqslant e^{-n^{\varepsilon}/16} = o(1).$$

(iii) We now prove $\mathbb{P}\left[\mathcal{E}_3 \mid \overline{\mathcal{E}_2}\right] < 1/4$. From Lemma 3.10.(*ii*) we get that w.p. $1/2$ a random walk starting at any node $u \in V_1$ does not leave G_1, *i. e.,* does not reach z^*, after $c_1 n$ time-steps for some constant $c_1 > 0$. It is easy to see that the number of visits to z^* required before the random walk hits G_2 instead of returning to G_1 is w.c.p. at least $\sqrt{\alpha'}/2$; this is because the fraction of edges from z^* to G_2 is $\sqrt{n/\alpha'}/(\sqrt{n/\alpha'}+\sqrt{n})$. Using a Chernoff bound, we conclude that any random walk starting at G_1 doesn't hit G_2 during the first $T = c_1 \cdot \sqrt{\alpha' n}/2$ time-steps with constant probability $p > 0$. Thus the probability that a random walk does not reach G_2 after $\lambda \cdot T$ time-steps is at least p^{λ}, for any integer $\lambda \geqslant 1$. Setting $\lambda = \varepsilon \cdot \log(1/p) \cdot \log(n/4)$, the probability that all of the at least $\frac{1}{4}n^{\varepsilon}$ random walks starting from G_1 reach G_2 within $\lambda \cdot T = \Omega(\sqrt{\alpha'} \cdot n \log n)$ steps is

$$\mathbb{P}\left[\mathcal{E}_3 \mid \overline{\mathcal{E}_2}\right] \leqslant (1-p^{\lambda})^{\frac{1}{4}n^{\varepsilon}} \leqslant 1/e.$$

□

The following lemma establishes a bound on the average hitting time.

Lemma 3.16. *Let G be the described at the beginning of Section 3.4. We have $t_{\text{avg-hit}} = \Omega(n^{3/2})$*

Proof. Consider a random walk that starts from an arbitrary vertex $u \in V$. By Lemma 3.10.(ii), every time a vertex z^i is visited, with probability at least $c > 0$ it takes $\Omega(n)$ time-steps to visit another vertex z^j, $j \neq i$. Using a Chernoff bound, it follows that with probability larger than $1/2$ it takes at least $\Omega(n^{3/2})$ time-steps to visit at least half of the nodes in $\{z^1, z^2, \ldots, z^{\kappa}\}$. By symmetry, it follows that for every vertex in a copy of G_1 there are $\Omega(n)$ vertices to which the hitting time is $\Omega(n^{3/2})$. Thus, by symmetry, $t_{\text{avg-hit}} = \sum_{u,v \in V} \pi(u) \cdot \pi(v) \cdot t_{\text{hit}}(u,v) = \Omega(n^2 \frac{1}{n^2} n^{3/2}) = \Omega(n^{3/2})$. □

4 Bounding $t_{\text{coal}} = O(t_{\text{hit}})$ for Almost-Regular Graphs

As mentioned in the introduction, the bound on t_{coal} in terms of t_{hit} will be based on the combination of two reduction results; first from $O(n)$ to $O(\log^3 n)$, and then from $O(\log^3 n)$ to $(\Delta/d)^{100}$; both taking $O(t_{\text{hit}})$ time each. In Section 4.1, we first develop concentration inequalties that will be needed for these reductions. Then in Section 4.2, we present the first and technically simpler reduction to $O(\log^3 n)$ walks, which is stated in Theorem 4.3. The proof basically combines the concentration inequalities with our well-known formula (1).

The derivation of the second reduction is done in Section 4.3. It is based on identifying nearly-regular and dense subsets S, which will contain enough vertices visited by a random walk, even if the walks only run for $o(t_{\text{hit}})$ steps (Lemma 4.5). The proof of this Lemma 4.5 also rests on the concentration inequalities we derive. The second reduction is then completed by Lemma 4.7, which uses the dense subsets S provided by Lemma 4.5 in order to prove that random walks are likely to collide. A more detailed proof outline can be found at the beginning of Section 4.3.

4.1 Concentration Inequalities for Random Walks

In this part, we derive several concentration inequalities for random walks that are new to the best of our knowledge. We point out that existing concentration inequalities tend to fail in our setting, since the events we are considering (like visits to a certain vertex or expected collisions with an unexposed walk) may only appear a small number of times during t_{hit} steps. Therefore, we have to develop new concentration inequalities that are parameterized by t_{hit}. Although the derivation is fairly elementary, the bounds are quite general and may complement existing bounds that are usually parameterized by the mixing time [10, 24]. In particular, our bounds are most useful when t_{mix} and t_{hit} are close, which is precisely the challenging regime for proving $t_{\text{coal}} = O(t_{\text{hit}})$. One limitation though is that our bounds only work for large deviations exceeding the expectation by a multiplicative factor.

Lemma 4.1. *Let $f : V \to [0,1]$ be any function over the vertices and $\overline{f} = \sum_{u \in V} f(u) \cdot \pi(u)$, and let T be any number of steps satisfying $T \geqslant t_{\text{hit}}$. Then for any random walk starting from an arbitrary vertex X_0,*

$$\mathbb{E}\left[\sum_{t=0}^{T} f(X_t)\right] \leqslant 8 \cdot T \cdot \overline{f}.$$

Furthermore, for any integer $\lambda \geqslant 1$,

$$\mathbb{P}\left[\sum_{t=0}^{T} f(X_t) \geqslant \lambda \cdot \left(16 \cdot T \cdot \overline{f} + 1\right)\right] \leqslant 2^{-\lambda}.$$

Moreover, suppose we have time-dependent functions, $f_t : V \to [0,1]$, $0 \leqslant t \leqslant T$, where T may be any integer (possibly smaller than t_{hit}.) Further assume that there is a universal bound $\Upsilon > 0$ so that for any $1 \leqslant s \leqslant T$ and $X_s = w$,

$$\mathbb{E}\left[\sum_{t=s}^{T} f_t(X_t) \mid X_s = w\right] \leqslant \Upsilon.$$

Then, again for any integer $\lambda \geqslant 1$,

$$\mathbb{P}\left[\sum_{t=0}^{T} f_t(X_t) \geqslant \lambda \cdot (2\Upsilon + 1)\right] \leqslant 2^{-\lambda}$$

Proof. We first prove that for all pairs of states $u, v \in V$ and any $T \geqslant t_{\text{hit}}$ that,

$$\sum_{t=0}^{T} p_{u,v}^{t} \leqslant 8 \cdot T \cdot \pi(v).$$

We will use $N_t(v)$ to denote the number of visits to v within t steps. Suppose for a sake of contradiction that $\sum_{t=0}^{T} p_{u,v}^{t} > 8 \cdot T \cdot \pi(v)$. Then, for an arbitrary vertex $w \in V$, by Markov's inequality, $\mathbb{P}[\,\tau(w,u) \leqslant 2\,t_{\text{hit}}\,] \geqslant \frac{1}{2}$. Therefore,

$$\mathbb{E}[\,N_{3T}(v) \mid X_0 = w\,] \geqslant \mathbb{P}[\,\tau(w,u) \leqslant 2\,t_{\text{hit}}\,] \cdot \mathbb{E}[\,N_T(v) \mid X_0 = u\,] \geqslant \frac{1}{2} \cdot \sum_{t=0}^{T} p_{u,v}^{t} > 4 \cdot T \cdot \pi(v).$$

Since this holds for every vertex $w \in V$, we also conclude $\mathbb{E}[\,N_{3T}(v) \mid X_0 \sim \pi\,] > 4T \cdot \pi(v)$. However, by definition of the stationary distribution, we also have $\mathbb{E}[\,N_{3T}(v) \mid X_0 \sim \pi\,] = 3T \cdot \pi(v)$, which yields the desired contradiction.

Now the first statement of the lemma follows simply by linearity of expectations:

$$\begin{aligned}\mathbb{E}\left[\sum_{t=0}^{T} f(X_t) \mid X_0 = w\right] &= \mathbb{E}\left[\sum_{t=0}^{T}\sum_{u\in V} \mathbf{1}_{X_t=u}\cdot f(u) \mid X_0 = w\right] \\ &= \sum_{u\in V}\sum_{t=0}^{T} f(u)\cdot \mathbb{E}\left[\mathbf{1}_{X_t=u} \mid X_0 = w\right] \\ &= \sum_{u\in V} f(u)\cdot \sum_{t=0}^{T} p^t_{w,u} \\ &\leqslant \sum_{u\in V} f(u)\cdot 8T\cdot \pi(u) = 8\cdot T\cdot \overline{f}.\end{aligned}$$

We now prove the second statement. By Markov's inequality,

$$\mathbb{P}\left[\sum_{t=0}^{T} f(X_t) \geqslant 16\cdot T\cdot \overline{f} \mid X_0 = w\right] \leqslant \frac{1}{2}.$$

Hence with $\tau := \min\left\{s\in\mathbb{N}\colon \sum_{t=0}^{s} f(X_t) \geqslant 16\cdot T\cdot \overline{f}\right\}$ we have,

$$\mathbb{P}\left[\tau \leqslant T \mid X_0 = w\right] \leqslant \frac{1}{2}.$$

Since f is bounded by 1, we get $\sum_{t=0}^{\tau} f(X_t) \leqslant 16\cdot T\cdot \overline{f} + 1$ and therefore,

$$\mathbb{P}\left[\sum_{t=0}^{T} f(X_t) \geqslant \lambda\cdot 16\cdot T\cdot \overline{f} + \lambda \mid X_0 = w\right] \leqslant \left(\max_{v\in V}\mathbb{P}\left[\tau \leqslant T \mid X_0 = v\right]\right)^{\lambda} \leqslant 2^{-\lambda}.$$

The third statement is derived in exactly the same way. □

The third statement of Lemma 4.1 is very useful in that it can be used the following concentration inequality on $\widetilde{Z}$. Notice that the variable random variable $\widetilde{Z}$ is defined using only one random walk $(X_t)_{t\geqslant 0}$, but it can be viewed as the expected number of collisions of random walk $(X_t)_{t\geqslant 0}$ with another (unexposed) random walk $(Y_t)_{t\geqslant 0}$, starting from the same vertex u.

Lemma 4.2. *Let S be any subset of vertices such that the degree of any pair of vertices in S differs by at most a factor of κ. Then consider any random walk $(X_t)_{t\geqslant 0}$ which starts an arbitrary vertex $u\in S$. Then for any $T\geqslant t_{\mathrm{hit}}$,*

$$\widetilde{Z} := \sum_{t=0}^{T} \mathbf{1}_{x_t\in S}\cdot p^t_{u,x_t},$$

we have that

$$\mathbb{E}\left[\widetilde{Z}\right] \leqslant 8\kappa\cdot T\cdot \max_{w\in S}\pi(w) =: \Upsilon.$$

Furthermore, for any $\lambda\geqslant 1$,

$$\mathbb{P}\left[\widetilde{Z} \geqslant \lambda\cdot(2\Upsilon + 1)\right] \leqslant 2^{-\lambda}.$$

Proof. First note that $\widetilde{Z}$ is a random variable over the walk $(X_0 = u, X_1, X_2, \ldots, X_T)$ with $u \in S$.

Let us first upper bound the expectation of $\widetilde{Z}$:

$$\mathbb{E}\left[\widetilde{Z}\right] \leqslant \sum_{t=0}^{T} \sum_{v \in S} p^t_{u,v} \cdot p^t_{u,v} \leqslant \kappa \cdot \sum_{t=0}^{T} \sum_{v \in S} p^t_{u,v} \cdot p^t_{v,u} \leqslant \kappa \cdot \sum_{t=0}^{T} p^{2t}_{u,u} \leqslant 8\kappa \cdot T \cdot \pi(u),$$

where the second inequality is due to reversibility, *i. e.*, $p^t_{u,s} \cdot \pi(u) = p^t_{s,u} \cdot \pi(s)$ and the fact that the degrees in S differ by a factor of at most κ, and the fourth inequality uses that $p^t_{u,u}$ is non-decreasing (Lemma A.1) and the first statement of Lemma 4.1.

Furthermore, suppose now that we condition on the walk $(X_t)_{t \geqslant 0}$ being at an arbitrary vertex $w \in S$ at step s, where $1 \leqslant s \leqslant T$. Then the contribution towards $\widetilde{Z}$ would be at most

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=s}^{T} \mathbf{1}_{X_t \in S} \cdot p^t_{u,x_t} \mid X_s = w\right] &= \sum_{t=s}^{T} \sum_{v \in S} p^{t-s}_{w,v} \cdot p^t_{u,v} \\
&\leqslant \kappa \cdot \sum_{t=s}^{T} \sum_{v \in S} p^{t-s}_{w,v} \cdot p^t_{v,u} \\
&\leqslant \kappa \cdot \sum_{t=s}^{T} p^{2t-s}_{u,w} \\
&\leqslant \kappa \cdot 8 \cdot T \cdot \pi(w) \leqslant \Upsilon,
\end{aligned}$$

where the penultimate inequality is due to the first statement of Lemma 4.1, applied to the number of visits to w of a random walk of length $T \geqslant t_{\text{hit}}$, i.e., $f(v) := \mathbf{1}_{v=w}$. Finally, by the third statement of Lemma 4.1, applied to the functions $f_t(v) := \mathbf{1}_{v \in S} \cdot p^t_{u,v}$, $0 \leqslant t \leqslant T$,

$$\mathbb{P}\left[\widetilde{Z} \geqslant \lambda \cdot \left(16\kappa \cdot T \cdot \max_{w \in S} \pi(w) + 1\right)\right] \leqslant 2^{-\lambda}.$$

□

4.2 Reducing the Walks from n to $O(\log^3 n)$ in $O(t_{\text{hit}})$

We now present our first reduction result that reduces the number of walks from n to $O(\log^3 n)$ in $O(t_{\text{hit}})$ time.

Theorem 4.3. *Let $G = (V, E)$ be an arbitrary, possibly non-regular, graph. Then after $O(t_{\text{hit}})$ steps, the number of walks can be reduced from n to $O(\log^3 n)$ with probability at least $1 - n^{-1}$.*

Thanks to $t_{\text{meet}} \leqslant 4t_{\text{hit}}$ and Proposition 3.3, this result implies, among other things, a bound of $t_{\text{coal}} = O(t_{\text{hit}} \cdot \log\log n)$ for any graph.

Proof of Theorem 4.3. First consider any random walk $(X_t)_{t=0}^{t_{\text{sep}}+2t_{\text{hit}}}$, that reaches an arbitrary vertex u at time t_{sep}. Next divide all vertices in V into buckets $S_i = \{v \in V \colon \deg(u) \in (2^{i-1}, 2^i]\}$. Using Lemma 4.1, it follows that any bucket S_i with $\pi(S_i) \leqslant 1/\log^3 n$ gets at most $t_{\text{hit}}/\log n$ visits within the time-interval $[t_{\text{sep}}, t_{\text{sep}} + 2t_{\text{hit}}]$ with probability at least $1 - n^{-2}/2$. Hence with

$$S := \bigcup_{\substack{1 \leqslant i \leqslant \log n \colon \\ \pi(S_i) \geqslant 1/\log^3 n}} S_i,$$

it follows that S gets at least $t_{\mathrm{hit}} - t_{\mathrm{hit}}/\log n \geqslant t_{\mathrm{hit}}/2$ visits with probability at least $1 - n^{-2}/2$.

Let us now consider any S_i with $\pi(S_i) \geqslant 1/(\log n)^3$, and define

$$\widetilde{Z}_i(s) := \sum_{t=s}^{t_{\mathrm{sep}}+2t_{\mathrm{hit}}} \mathbf{1}_{x_t \in S_i} \cdot p^{t-s}_{x_s,x_t},$$

where s is any time-step satisfying $t_{\mathrm{sep}} \leqslant s \leqslant t_{\mathrm{sep}}+2t_{\mathrm{hit}}$. Notice that by Lemma 4.2, with $\lambda \geqslant c\cdot\log n$ for a proper constant $c > 0$, we have

$$\mathbb{P}\left[\widetilde{Z}_i(s) \geqslant c\log n \cdot 10 \cdot t_{\mathrm{hit}} \cdot \max_{w\in S_i}\pi(w)\right] \leqslant n^{-8},$$

having used the fact that $t_{\mathrm{hit}} \geqslant 1/\min_{v\in V}\pi(v)$. Since $t_{\mathrm{hit}} = O(n^3)$ and there are at most $\log n$ buckets, by the Union bound,

$$\mathbb{P}\left[\bigcup_{t_{\mathrm{sep}}\leqslant s\leqslant t_{\mathrm{sep}}+2t_{\mathrm{hit}}}\left\{\widetilde{Z}_i(s) \geqslant c\log n \cdot 10 \cdot t_{\mathrm{hit}} \cdot \max_{w\in S_i}\pi(w)\right\}\right] \leqslant n^{-2}/2.$$

Hence by the Union bound, with probability at least $1-n^{-2}$, the trajectory $(x_0, x_1, \ldots, x_{t_{\mathrm{sep}}+2t_{\mathrm{hit}}})$ of $(X_t)_{t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+2t_{\mathrm{hit}}}$ is *good*, i.e., its trajectory (i) makes at least $t_{\mathrm{hit}}/2$ visits to S and (ii) all $\widetilde{Z}_i(s)$ are bounded by $c\log n\cdot 10\cdot t_{\mathrm{hit}}\cdot\max_{w\in S_i}\pi(w)$. In the following, condition on $(X_t)_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+2t_{\mathrm{hit}}}$ being a good random walk, and let us denote by $(x_0, x_1, \ldots, x_{t_{\mathrm{sep}}+2t_{\mathrm{hit}}})$ the deterministic trajectory.

Since S gets at most $t_{\mathrm{hit}}/2$ visits in the time-interval $[t_{\mathrm{sep}}, t_{\mathrm{sep}}+2t_{\mathrm{hit}}]$, by the pigeonhole principle, there must be at least one bucket S_j with $\pi(S_j) \geqslant 1/\log^3 n$ so that bucket S_j gets at least $(1/8)\cdot t_{\mathrm{hit}}/\log n$ visits in that time-interval. We shall now prove that any other random walk $(Y_t)_{t=0}^{t_{\mathrm{sep}}+2t_{\mathrm{hit}}}$, starting from an arbitrary vertex $w \in V$, collides with this deterministic trajectory on a vertex in S_j in the time-interval $[t_{\mathrm{sep}}, t_{\mathrm{sep}} + 2t_{\mathrm{hit}}]$ with probability at least $\Omega(1/\log^2 n)$. To this end, let us define

$$Z := \sum_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+2t_{\mathrm{hit}}} \mathbf{1}_{X_t\in S_j}\cdot\mathbf{1}_{X_t=Y_t}.$$

Since for any $t \geqslant t_{\mathrm{sep}}$ we have $p^t_{w,v} \geqslant \frac{1}{2}\cdot\pi(v)$ for any pair of vertices $w, v \in V$, it follows that

$$\begin{aligned}
\mathbb{E}\left[\, Z \mid (x_0, x_1, \ldots, x_{t_{\mathrm{sep}}+2t_{\mathrm{hit}}})\,\right] &\geqslant \sum_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+2t_{\mathrm{hit}}}\sum_{v\in S_j}\mathbf{1}_{x_t=v}\cdot p^t_{w,v} \\
&\geqslant \sum_{v\in S_j}\sum_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+2t_{\mathrm{hit}}}\mathbf{1}_{x_t=v}\cdot\frac{1}{2}\cdot\pi(v) \\
&\geqslant \frac{1}{2}\cdot\min_{w\in S_j}\pi(w)\cdot\sum_{v\in S_j}\sum_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+2t_{\mathrm{hit}}}\mathbf{1}_{x_t=v} \\
&\geqslant \frac{1}{16}\cdot\min_{w\in S_j}\pi(w)\cdot\frac{1}{\log n}\cdot t_{\mathrm{hit}},
\end{aligned}$$

where the last inequality holds because the deterministic path $(x_0, x_1, \ldots, x_{t_{\mathrm{sep}}+2t_{\mathrm{hit}}})$ makes at least $(1/8)\cdot t_{\mathrm{hit}}/\log n$ visits to S_j.

Furthermore, since the deterministic walk $(x_0, x_1, \ldots, x_{t_{\text{sep}}+2t_{\text{hit}}})$ satisfies invariant (ii), conditional on Y_t having its first collision with X_t at step s on a vertex $x_s \in S_j$,

$$\begin{aligned}\mathbb{E}\left[Z \mid (x_0, x_1, \ldots, x_{t_{\text{sep}}+2t_{\text{hit}}}), Z \geqslant 1 \right] &\leqslant \max_{t_{\text{sep}} \leqslant s \leqslant t_{\text{sep}}+2t_{\text{hit}}} \sum_{t=s}^{t_{\text{sep}}+2t_{\text{hit}}} \mathbf{1}_{x_t \in S_j} \cdot p^{t-s}_{x_s, x_t} \\ &= O(\log n \cdot t_{\text{hit}} \cdot \max_{w \in S_j} \pi(w)).\end{aligned}$$

Combining our last two bounds yields

$$\begin{aligned}\mathbb{P}\left[Z \geqslant 1 \mid (x_0, x_1, \ldots, x_{t_{\text{sep}}+2t_{\text{hit}}}) \right] &= \frac{\mathbb{E}\left[Z \mid (x_0, x_1, \ldots, x_{t_{\text{sep}}+2t_{\text{hit}}}) \right]}{\mathbb{E}\left[Z \mid (x_0, x_1, \ldots, x_{t_{\text{sep}}+2t_{\text{hit}}}), Z \geqslant 1 \right]} \\ &= \Omega(1/\log^2 n) =: p \end{aligned} \tag{12}$$

To complete the proof of the theorem, divide the $k > 2 \cdot 1/p \cdot \log n$ random walk arbitrarily into two disjoint groups G_1 and G_2 such that $|G_1| = 1/p \cdot \log n$. By the above argument, any random walk $(X_t)_{t=0}^{t_{\text{sep}}+2t_{\text{hit}}}$ in G_2 will be good with probability at least $1 - n^{-2}$. Hence by Markov's inequality, all trajectories of the random walks in G_2 are good with probability at least $1 - n^{-1}$. Conditioning on this event, (12) shows that any from the random walks in G_1 collides with the trajectory of any fixed good random walk in G_2 in the time-interval $[t_{\text{sep}}, t_{\text{sep}}+2t_{\text{hit}}]$ with probability at least p. By independence of these events across random walks in G_1, random walk $(X_t)_{t=0}^{t_{\text{sep}}+2t_{\text{hit}}}$ is not eliminated with probability at most

$$(1-p)^{2\frac{1}{p} \cdot \log n} \leqslant n^{-2}.$$

Combining everything, it follows that with probability at least $1 - 2 \cdot n^{-1}$ all random walks in G_2 are eliminated. The claim follows by noting that $t_{\text{sep}} + 2t_{\text{hit}} = O(t_{\text{mix}} + t_{\text{hit}}) = O(t_{\text{hit}})$. □

4.3 Reducing the Walks from $\log^4 n$ to $(\Delta/d)^{O(1)}$ in $O(t_{\text{hit}})$

Theorem 4.4. *Let $G = (V, E)$ be any graph with maximum degree Δ and average degree d. Then the expected time to reduce the number of walks from $\log^4 n$ to $(\Delta/d)^{100}$ is at most $O(t_{\text{hit}})$.*

In comparison with Theorem 4.3, the reduction in Theorem 4.4 is more subtle, as there might be a sub-logarithmic number of random walks preventing us from using the simple bucketing-argument into "nearly-regular" partitions used in Theorem 4.3. Furthermore, Theorem 4.3 achieved the reduction in just a single phase of $\Theta(t_{\text{hit}})$ steps in order to "mix" all random walks. Here however, we are only able to prove an "exponential" progress and need to cope with random walks whose distribution may be far from the stationary distribution. Specifically, if there are k random walks left, we will analyze a phase of length t_{hit}/κ, with $k = \kappa^{100}$, and show that a constant fraction of random walks will be eliminated. To account for the fact that the random walks are not mixed, we will identify certain "dense" subsets S having two crucial properties. First, these subsets have a stationary mass very close to 1, and the distance decreases in κ. Secondly, these subsets consist only of those vertices for which the stationary probability is close to $1/n$.

While the existence of these sets S can be shown by a series of counting arguments, more involved methods are needed to show that random walks will pay enough visits to vertices in S within t_{hit}/κ steps. Here, we apply our new concentration inequalities (Lemma 4.1) to rule out events where the random walk spends too many steps in the set $V \setminus S$. This is formally shown in

Lemma 4.5, establishing that for any random walk i, with large probability, there exists a large subset $S(i) \subseteq S$ of vertices that will all be visited "sufficiently" often.

It turns out that through a careful choice of κ, we end up in a favorable situation where for each walk $1 \leqslant i \leqslant k$, we have a subset $S(i) \subseteq S$ with $|S(i)| \geqslant n/\kappa^{\Omega(1)}$. Thanks to the large gap between k and κ, it turns out that an overwhelmingly large fraction of these subsets $S(i)$, $1 \leqslant i \leqslant k$, must overlap. Unfortunately, we still need to make sure that random walks will visit the right vertices at the right time. To ensure this, we will eliminate "highly surprising visits", which are visits to vertices at a time when the probability for this to happen is at most $\kappa^{-O(1)}$ (Lemma 4.5). Put differently, we show that an overwhelming fraction of visits occurs indeed at time-steps for which the probability is at least $\kappa^{-O(1)}$.

Equipped with these subsets $S(i)$, we regard the "not too surprising" visits as a balls-into-bins configuration, where each ball on a bin (vertex) is associated to a walk i which may visit this vertex. Through a series of counting arguments, we establish that most random walk make mostly visits to vertices at times, when a lot of other random walks could also visit that vertex (we refer to Figure 4 for an illustration).

After all these preparations, we are able to analyze the coalescing process and achieve the desired reduction in the number of the random walks. Similarly to previous analyses, we use a division of random walks into groups G_1 and G_2. The roles of G_1 and G_2 are as before; walks in G_1 are merely used to eliminate walks in G_2. This time, however, this division is completely uniformly at random. We establish that for most fixed random walks $1 \leqslant i \leqslant k$, conditional on being in group G_2, there is a constant probability of picking a trajectory that will likely lead to an intersection with any of the other $k-1$ unexposed random walk. This second part of the proof is formalized in Lemma 4.7.

Combining the two steps of the proof, the structural result in Lemma 4.5 with the probabilistic analysis in Lemma 4.7, it immediately follows that the number of walks can be reduced by a constant factor within $O(t_{\text{hit}}/\kappa)$ steps, yielding Theorem 4.4. For our analysis, a necessary condition is that $k = \kappa^{100}$ is sufficiently large compared to Δ/d; this is crucial for the existence of the dense subset S, for which we can analyze the number of collisions in Lemma 4.7.

Before giving the formal proof of Theorem 4.8, we introduce additional notation. Recall that k is the number of random walks at a certain time, w.l.o.g., say $t = 0$. Consider a fixed random walk $(X_t)_{t=0}^{\infty}$ with label $1 \leqslant i \leqslant k$, where $k = \kappa^{100}$ that is run for

$$\tau := 4t_{\text{hit}}/\kappa$$

steps, and starts at an arbitrary vertex $X_0 = u_i$. Let

$$S_0 = S_0(i) := \{v \in V \colon \mathbb{P}\left[\,\tau(u_i, v) \leqslant \tau\,\right] \geqslant \kappa^{-21}\}.$$

Furthermore, let

$$t(v) := \min\left\{t \colon \mathbb{P}\left[\,\tau(u_i, v) \leqslant t\,\right] \geqslant \kappa^{-21}\right\}.$$

Let $N_t(v)$ denote the number of visits of the random walk to v within t time-steps and let $S_1(i)$ be the set of vertices v that get at least $\tau/\kappa^4 \cdot \pi(v)$ visits between the time steps $t(v)$ and τ; in symbols

$$S_1 = S_1(i) := \left\{v \in V \colon N_\tau(v) - N_{t(v)}(v) \geqslant \tau/\kappa^4 \cdot \pi(v)\right\}.$$

It is easy to verify that we always have $S_1 \subseteq S_0$.

We now to state our structural lemma, providing lower bounds on the size of $S_1(i)$. For technical reasons, we even allow the restriction of $S_1(i)$ to any given "dense" set S that has a large stationary mass and contains only vertices with sufficiently high degree. This "projection" is not required on

regular graphs, where we could simply work with all vertices, *i. e.*, $S = V$. However, for non-regular graphs, the projection on S is very useful since on the set S, the random walk behaves essentially like a random walk on a regular graph.

Lemma 4.5. *Let $G = (V, E)$ be an arbitrary graph and let $\kappa \geqslant 20$ be any integer. Let $S \subseteq V$ be such that $\pi(S) \geqslant 1 - 1/\kappa^2$ and $1/(\kappa^3 n) \leqslant \pi(u) \leqslant \kappa^3/n$ for every $u \in S$. Consider any random walk with label i. Then, we have that $|S_1(i) \cap S| \geqslant n/\kappa^8$ w.p. at least $1 - 6/\kappa$.*

In other words, the lemma tells us that a significant fraction of visits of random walk i will be to vertices that are not "too surprising" to be visited.

Proof of Lemma 4.5. First we bound the number of visits to $V \setminus S$. To this end, let

$$\widetilde{C} := \sum_{t=0}^{\tau} \sum_{v \notin S} \mathbf{1}_{X_t = v}.$$

By Lemma 4.1,

$$\mathbb{E}\left[\widetilde{C}\right] \leqslant 8 t_{\text{hit}} \cdot \pi(V \setminus S) \leqslant 8 t_{\text{hit}}/\kappa^2 = 2\tau/\kappa.$$

Hence by Markov's inequality,

$$\mathbb{P}\left[\widetilde{C} \geqslant \tau/2\right] \leqslant \mathbb{P}\left[\widetilde{C} \geqslant \kappa/4 \cdot \mathbb{E}[C]\right] \leqslant \frac{4}{\kappa}.$$

Next for any fixed vertex $v \in V$, we know that for any $\lambda \geqslant 1$ the probability that the random walk makes more than $32\lambda \cdot t_{\text{hit}} \cdot \pi(v)$ visits to v is at most $2^{-\lambda}$, by Lemma 4.1 (here we are using the simple fact that $t_{\text{hit}} \cdot \pi(v) \geqslant 1$). Recall that $N_t(v)$ denotes the number of visits of the random walk to v within t time-steps. Define

$$B := \sum_{v \in V} \mathbf{1}_{N_\tau(v) \geqslant \kappa^2 \cdot t_{\text{hit}} \cdot \pi(v)} \cdot N_\tau(v).$$

Then

$$\begin{aligned}
\mathbb{E}[B] &= \sum_{v \in V} \sum_{j = \kappa^2 \cdot t_{\text{hit}} \cdot \pi(v)}^{\infty} j \cdot \mathbb{P}[N_\tau(v) = j] \\
&\leqslant \sum_{v \in V} \sum_{\sigma = \kappa^2}^{\infty} (\sigma + 1) \cdot t_{\text{hit}} \cdot \pi(v) \cdot \mathbb{P}[N_\tau(v) \geqslant \sigma \cdot t_{\text{hit}} \cdot \pi(v)] \\
&\leqslant \sum_{v \in V} t_{\text{hit}} \cdot \pi(v) \cdot \sum_{\sigma = \kappa^2}^{\infty} 2\sigma \cdot 2^{-\sigma/32} \\
&\leqslant 2 t_{\text{hit}} \cdot \kappa^{-10} \leqslant \tau \cdot \kappa^{-9},
\end{aligned}$$

where we used that $\kappa \geqslant 20$ and $\sum_{v \in V} \pi(v) = 1$. Hence, by Markov's inequality,

$$\mathbb{P}\left[B \geqslant \tau \cdot \kappa^{-8}\right] \leqslant \frac{1}{\kappa}.$$

Suppose now that $\left\{\widetilde{C} \leqslant \tau/2\right\}$ and $\left\{B \leqslant \tau \cdot \kappa^{-8}\right\}$ both occur. Next define

$$\widetilde{S} := \left\{u \in S \colon N_\tau(u) \geqslant \pi(u) \cdot \tau/\kappa^4\right\},$$

i. e., $\widetilde{S}$ are all vertices $v \in S$ that are visited at least $\tau/\kappa^4 \cdot \pi(v)$ times before time-step τ. Notice that $\widetilde{S}$ is a (random) set that depends on the realization of the walk. Conditioning on $B \leqslant \tau/\kappa^8$ and $\widetilde{C} \leqslant \tau/2$, we will show by pigeonhole principle $|\widetilde{S}| \geqslant 2n/\kappa^8$: Suppose for the sake of contradiction that $|\widetilde{S}| < 2n/\kappa^8$. Then, using $\pi_{\max} = \max_{u \in V} \pi(u) \leqslant \kappa/n$, we have that the total number of visits to nodes in S is at most

$$\begin{aligned}
\sum_{u \in S \setminus \widetilde{S}} N_\tau(u) + \sum_{u \in \widetilde{S}} N_\tau(u) &\leqslant \sum_{u \in S \setminus \widetilde{S}} \pi(u) \cdot \tau/\kappa^4 + \sum_{u \in \widetilde{S} \colon N_\tau(u) \leqslant \kappa^2 \cdot t_{\text{hit}} \cdot \pi(u)} N_\tau(u) + \sum_{u \in \widetilde{S} \colon N_\tau(u) > \kappa^2 \cdot t_{\text{hit}} \cdot \pi(u)} N_\tau(u) \\
&\leqslant (|S| - |\widetilde{S}|) \cdot (\kappa^3/n) \cdot \tau/\kappa^4 + \left(|\widetilde{S}| \cdot \max_{u \in S} \pi(u) \cdot t_{\text{hit}} \cdot \kappa^2 + B \right) \\
&< \tau/\kappa + 2n/\kappa^8 \cdot (\kappa^3/n) \cdot \tau \cdot \kappa \cdot \kappa^2 + \tau/\kappa^8 \\
&\leqslant \frac{1}{2}\tau,
\end{aligned}$$

a contradiction to the event $\left\{\widetilde{C} \leqslant \tau/2\right\}$.

Hence with probability $1 - \mathbb{P}\left[\left\{\widetilde{C} \geqslant \tau/2\right\} \cup \{B \geqslant \kappa \cdot \mathbb{E}[B]\} \right] \geqslant 1 - \frac{5}{\kappa}$ there is a set $\widetilde{S}$ with $|\widetilde{S}| \geqslant 2n/\kappa^8$.

Finally, we will upper bound the number of "surprising" visits, which are visits to vertices that happen too early. That is, we will upper bound the number of visits to vertices v before time $t(v)$; in symbols,

$$\widetilde{B} := \sum_{v \in V} \sum_{0 \leqslant t < t(v)} \mathbf{1}_{X_t = v}.$$

By linearity of expectations, and for every v, the number of visits to v is maximized when starting at v,

$$\mathbb{E}\left[\widetilde{B}\right] \leqslant \sum_{v \in V} \left(\kappa^{-21} \cdot \sum_{t=0}^{\tau} p_{v,v}^t \right),$$

since with probability at most κ^{-21}, the random walk makes a visit to v before step $t(v)$, and after that, the expected number of returns to v is at most $\sum_{t=0}^{\tau} p_{v,v}^t$. Since $\sum_{t=0}^{\tau} p_{v,v}^t \leqslant 32 t_{\text{hit}} \cdot \pi(v) = 8\kappa^3 \cdot \tau \cdot \pi(v)$ by Lemma 4.1, it follows that

$$\mathbb{E}\left[\widetilde{B}\right] \leqslant \kappa^{-21} \cdot 8\kappa^3 \tau \cdot \sum_{v \in V} \pi(v) \leqslant \tau/2 \cdot \kappa^{-16},$$

and thus by Markov's inequality,

$$\mathbb{P}\left[\widetilde{B} \geqslant \tau/2 \cdot \kappa^{-15}\right] \leqslant \mathbb{P}\left[\widetilde{B} \geqslant \kappa \cdot \mathbb{E}\left[\widetilde{B}\right]\right] \leqslant \frac{1}{\kappa}.$$

Hence, the total number of visits to vertices in V before time $t(v)$ is at most $\tau/2 \cdot \kappa^{-15}$ with probability at least $1 - 1/\kappa$.

By union bound, it follows that with probability at least $1 - 6/\kappa$ there is a subset $S_1 = S_1(i) \subseteq \widetilde{S} \subseteq S$ such that each vertex $v \in S_1$ is visited at least $(1/2)\pi(v) \cdot \tau/\kappa^4$ times between time steps $t(v)$ and τ with

$$|S_1| \geqslant |\widetilde{S}| - \frac{\widetilde{B}}{(1/2)\min_{w \in S} \pi(w) \cdot \tau/\kappa^4} \geqslant \frac{2n}{\kappa^8} - \frac{(1/2)\tau \cdot \kappa^{-15}}{(1/2)1/(\kappa^3 n) \cdot \tau/\kappa^4} = \frac{n}{\kappa^8}.$$

The second part holds trivially thanks to the definitions of S_1 and S_0. □

Before turning to the proof of Lemma 4.7, we introduce another piece of notation. For any random walk $1 \leqslant i \leqslant k$ denoted by $(X_t)_{t \geqslant 0}$, we call a time-step $s \in [0, \tau]$ is *bad* if $x_t \in S$ and additionally,

$$\sum_{t=s}^{\tau} \mathbf{1}_{x_t \in S} \cdot p^{t-s}_{x_s, x_t} \geqslant 16 \cdot \kappa^{10} \cdot t_{\text{hit}}/n.$$

Intuitively a time-step t is bad, if the expected number of collisions for another random walk starting at vertex $x_t \in S$ at step t with the walk $(x_t, x_{t+1}, \ldots)$ is too large.

Lemma 4.6. *Consider a random walk $(X_t)_{t=0}^{\tau}$ of length $\tau = 4t_{\text{hit}}/\kappa$. Let $S \subseteq V$ be such that $\pi(S) \geqslant 1 - 1/\kappa^2$ and $1/(\kappa^3 n) \leqslant \pi(u) \leqslant \kappa^3/n$ for every $u \in S$. Then with probability at least $1 - 1/\kappa$, there are at most $\tau \cdot 2^{-\kappa}$ bad time-steps $t \in [0, \tau]$.*

Proof. First, let us fix any step $s \in [0, \tau]$, and following the notation of Lemma 4.2, let

$$\widetilde{Z}(s) := \sum_{t=s}^{t_{\text{hit}}} \mathbf{1}_{x_t \in S} \cdot p^{t-s}_{x_s, x_t}.$$

Then, by Lemma 4.2, for any vertex $u \in S$,

$$\mathbb{E}\left[\widetilde{Z}(s) \mid X_s = u \right] \leqslant 8\kappa^6 \cdot t_{\text{hit}} \cdot \max_{w \in S} \pi(w) \leqslant 8\kappa^9 \cdot t_{\text{hit}}/n =: \Upsilon.$$

Since $\Upsilon \geqslant 2$, the concentration inequality in Lemma 4.2 implies

$$\mathbb{P}\left[\widetilde{Z}(s) \geqslant 16 \cdot \kappa^{10} \cdot t_{\text{hit}}/n \right] \leqslant \mathbb{P}\left[\widetilde{Z}(s) \geqslant 2\kappa \cdot (2\Upsilon + 1) \right] \leqslant 2^{-2\kappa}.$$

Now let B denote the number of bad time-steps, i.e.,

$$B := \left| \left\{ 0 \leqslant s \leqslant \tau \colon \widetilde{Z}(s) \geqslant 4 \cdot \kappa^{10} \cdot t_{\text{hit}}/n \right\} \right|.$$

Then,

$$\mathbb{E}[B] \leqslant \tau \cdot 2^{-2\kappa},$$

and a simple application of Markov's inequality implies the claim. ∎

Using the sets $S_0(i)$ and $S_1(i)$ from Lemma 4.5, we can now analyze the collisions restricted to these sets:

Lemma 4.7. *Let $G = (V, E)$ be an arbitrary graph and suppose there are $k = \kappa^{100}$ random walks remaining, where $\kappa \geqslant 20$. Further suppose that each random walk with label $1 \leqslant i \leqslant k$ satisfies $\mathbb{P}\left[|S_1(i)| \geqslant n/\kappa^8 \right] \geqslant 1 - 6/\kappa$, and all vertices $u \in S_1(i)$ satisfy $1/(\kappa^3 n) \leqslant \pi(u) \leqslant \kappa^3/n$. Then with constant probability > 0, a constant fraction of the k walks is eliminated after $\tau = 4t_{\text{hit}}/\kappa$ steps.*

Proof. Consider any of the $k = \kappa^{100}$ random walks with label $1 \leqslant i \leqslant \kappa^{100}$, denoted by $(X^i_t)_{t=0}^{\tau}$. The random walk may start from an arbitrary vertex $X^i_0 = u_i$ and is run for $\tau = 4t_{\text{hit}}/\kappa$ steps.

We now consider the following balls-into-bins configuration. Every vertex in V corresponds to a bin. For every walk j and $u \in S_0(j)$, we place a ball with label $(j, t(j,u))$ into bin u. We call a ball with label $(j, t(j,u))$ in bin u *bad* if there are less than κ^{55} other balls i in the same bin such that either (i) $t(i,u) < t(j,u)$ or (ii) $t(i,u) = t(j,u)$ and $i < j$. Next define a random walk j to be

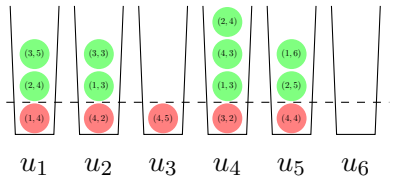

i	$t(i,u_1)$	$t(i,u_2)$	$t(i,u_3)$	$t(i,u_4)$	$t(i,u_5)$	$t(i,u_6)$
1	4	3	$>\tau$	3	6	$>\tau$
2	4	$>\tau$	$>\tau$	4	5	$>\tau$
3	5	3	$>\tau$	2	$>\tau$	$>\tau$
4	$>\tau$	2	5	3	4	$>\tau$

Figure 4: Illustration of the balls-into-bins configuration of $k = 4$ walks labeled $1, 2, 3, 4$ into 6 vertices $u_1, u_2, \ldots, u_6$. In the illustration, a ball $(i, t(i,u))$ on bin u is good if there is at least one other ball with $t(i,u) < t(j,u)$ (or $t(i,u) = t(j,u)$ and $i < j$). For each of the random walks 1, 2 and 3, there are at least two good balls, while for random walk 4 only one ball is good.

bad if at least n/κ^9 bad balls have label j and otherwise we call j *good*. Since there are at most $\kappa^{55} \cdot n$ bad balls, it follows that the number of *bad* walks $j \in \{1, \ldots, k\}$ is at most

$$\frac{\kappa^{55} \cdot n}{n/\kappa^9} \leqslant \kappa^{64}.$$

In the following, we will focus on the $\kappa^{100} - \kappa^{64}$ good walks and show that each such walk gets eliminated with constant probability > 0. Recall that for any fixed good walk i, we have at most n/κ^9 bad balls.

For each vertex appearing in the random set $S_1(i)$, we have placed one ball in the above balls-into-bins configuration thanks to $S_1(i) \subseteq S_0(i)$. Thus, with probability at least $1 - 6/\kappa - 1/\kappa = 1 - 7/\kappa$, at least $n/\kappa^8 - n/\kappa^9 \geqslant n/\kappa^9$ balls correspond to a vertex in $S_1(i)$ and are good. Let us denote by $S_2(i)$ this subset of $S_1(i)$ restricted to good balls; so $|S_2(i)| \geqslant n/\kappa^9$.

Furthermore, by Lemma 4.6, with probability at least $1 - 1/\kappa$, random walk i has at most $\tau \cdot 2^{-\kappa} \leqslant t_{\text{hit}}/\kappa^{20}$ bad time-steps $t \in [0, \tau]$, where we recall t is *bad* if

$$\sum_{t=s}^{\tau} \mathbf{1}_{x_t \in S} \cdot p^{t-s}_{x_s, x_t} \geqslant 16 \cdot \kappa^{10} \cdot t_{\text{hit}}/n.$$

Further recall that each vertex in $v \in S_2(i)$ is visited at least $\tau/\kappa^4 \cdot \pi(v) \geqslant \tau/(\kappa^7 n)$ times. Let $S_3(i) \subseteq S_2(i)$ denote the set of vertices $v \in S_2(i)$ such that at least $\tau/(\kappa^9 n)$ visits happen at a good time-step. We claim that $|S_3(i)| \geqslant n/\kappa^{10}$. Assume, for the sake of contradiction, that $|S_3(i)| < \frac{1}{2}|S_2(i)|$. That means, for at least $|S_2(i)| - |S_3(i)| \geqslant \frac{1}{2}|S_2(i)|$ vertices, we have at least $\tau/(\kappa^7 n) - \tau/(\kappa^9 n) \geqslant (1/2) \cdot \tau/(n\kappa^9)$ visits on a bad time step, which means that the total number of bad time steps would be at least

$$\frac{1}{2}|S_2(i)| \cdot \frac{1}{2} \cdot \frac{\tau}{n\kappa^9} \geqslant \frac{1}{4} \cdot \frac{\tau}{\kappa^{18}},$$

which contradicts the assumption that there are at most $t_{\text{hit}}/\kappa^{20}$ bad time steps in total.

Next we assign each random walk $1 \leqslant i \leqslant k = \kappa^{100}$ into group G_1 and G_2 independently and uniformly at random. Further, suppose that the trajectory of a walk i denoted by $(x_0 = u_i, x_1, \ldots, x_\tau)$ satisfies $|S_3(i)| \geqslant n/\kappa^{10}$ (which happens with constant probability $1 - 2/\kappa > 0$), and that the random walk i is assigned to the group G_2 (which happens with probability $1/2$). We will prove that conditional on the two above events, the random walk i will be eliminated by one of the walks in G_1 with at least constant probability.

To this end, consider now another random walk $j \neq i$ starting from an arbitrary vertex u_j. Denote by the set $S_4(i,j) \subseteq S_3(i)$ all vertices $u \in S_3(i)$ that random walk i visits u at least $\tau/(\kappa^9 n)$ times during good time-steps, and additionally, $t(j,u) \leqslant t(i,u)$. With reference to Figure 4, $S_4(i,j)$ counts the number of green balls corresponding to i that lie in a bin on top of a ball corresponding to j. By construction, for all vertices $u \in S_3(i)$ we have at least κ^{55} other random walks j such that $t(j,u) < t(i,u)$ (or $t(j,u) = t(i,u)$ and $j < i$). Hence, looking at all bins, the total number of balls which are below balls of walk i, we get, by summing up all the sizes of $S_4(i,j)$,

$$\sum_{j=1, j\neq i}^{\kappa^{100}} |S_4(i,j)| \geqslant \kappa^{55} \cdot |S_3(i)| \geqslant \kappa^{55} \cdot n/\kappa^9 = n\kappa^{46}. \tag{13}$$

Finally, we are now in a position to apply our common analysis method. Let $Z(i,j)$ denote the number of collisions between random walk j and i that happens on a vertex in $S_4(i,j)$ at a good time step, in symbols,

$$Z(i,j) := \sum_{1\leqslant t\leqslant \tau\colon t \text{ is good}} \mathbf{1}_{x_t \in S_4(i,j)} \cdot \mathbf{1}_{Y_t = x_t}.$$

With probability $1/2$, random walk j is in group G_1, and therefore

$$\mathbb{E}\left[\, Z(i,j) \mid (x_0, x_1, \ldots, x_\tau), i \in G_2 \,\right] = \frac{1}{2} \cdot \sum_{u \in S_4(i,j)} \sum_{\substack{1\leqslant t\leqslant \tau\colon \\ t \text{ is good and } x_t = u}} p^t_{u_j,u}.$$

We now would like to derive a lower bound on $p^t_{u_j,u}$, where $t \leqslant \tau$ is a good time-step with $X^i_t = u$, $u \in S_4(i,j)$. Since $S_4(i,j) \subseteq S_1(i)$, we have $t(j,u) \leqslant t(i,u) \leqslant \min\{s\colon x_s = u\} \leqslant t$, and therefore by definition of $t(j,u)$,

$$\begin{aligned}
p^t_{u_j,u} &= \sum_{s=0}^{t} \mathbb{P}\left[\, \tau(u_j,u) = s \,\right] \cdot p^{t-s}_{u,u} \\
&\geqslant \pi(u) \cdot \sum_{s=0}^{t} \mathbb{P}\left[\, \tau(u_j,u) = s \,\right] \\
&\geqslant 1/(\kappa^3 \cdot n) \cdot \mathbb{P}\left[\, \tau(u_j,u) \leqslant t \,\right] \\
&\geqslant 1/(\kappa^3 \cdot n) \cdot \mathbb{P}\left[\, \tau(u_j,u) \leqslant t(j,u) \,\right] \geqslant 1/(\kappa^{24} \cdot n).
\end{aligned}$$

Applying this and recalling that each vertex in $S_4(i,j)$ is visited at least $\tau/(\kappa^9 n)$ times during a good time step,

$$\begin{aligned}
\mathbb{E}\left[\, Z(i,j) \mid (x_0, x_1, \ldots, x_\tau), i \in G_2 \,\right] &\geqslant 2 \cdot |S_4(i,j)| \cdot \frac{\tau}{\kappa^9 \cdot n} \cdot \frac{1}{\kappa^{24} \cdot n} \\
&= 2 \cdot |S_4(i,j)| \cdot \frac{\tau}{\kappa^{33}} \cdot \frac{1}{n^2}.
\end{aligned}$$

Furthermore, by definition of a good time step,

$$\mathbb{E}\left[\, Z(i,j) \mid (x_0, x_1, \ldots, x_\tau), i \in G_2, Z(i,j) \geqslant 1 \,\right] \leqslant 16\kappa^{10} \cdot \frac{t_{\text{hit}}}{n} \leqslant \kappa^{12} \cdot \frac{\tau}{n}$$

Combining the last two inequalities yields

$$\begin{aligned}\mathbb{P}\left[\, Z(i,j) \geqslant 1 \mid (x_0, x_1, \ldots, x_\tau), i \in G_2 \,\right] &= \frac{\mathbb{E}\left[\, Z(i,j) \mid (x_0, x_1, \ldots, x_\tau), i \in G_2 \,\right]}{\mathbb{E}\left[\, Z(i,j) \mid (x_0, x_1, \ldots, x_\tau), i \in G_2, Z \geqslant 1 \,\right]} \\ &\geqslant \frac{|S_4(i,j)|}{n \cdot \kappa^{45}}\end{aligned}$$

Recalling that by (13), $\sum_{j=1, j\neq i}^{\kappa^{100}} |S_4(i,j)| \geqslant n\kappa^{46}$, we conclude that

$$\begin{aligned}\mathbb{P}\left[\cup_{j=1,j\neq i}^{\kappa^{100}} \{Z(i,j) \geqslant 1\} \mid (x_0, x_1, \ldots, x_\tau), i \in G_2 \right] &\geqslant 1 - \prod_{j=1, j\neq i}^{\kappa^{100}} \left(1 - \frac{|S_4(i,j)|}{n \cdot \kappa^{45}}\right) \\ &\geqslant 1 - \exp\left(-\frac{1}{n \cdot \kappa^{45}} \cdot \sum_{j=1, j\neq i}^{\kappa^{100}} |S_4(i,j)|\right) \\ &\geqslant 1 - \exp(-\kappa) \geqslant 1/2.\end{aligned}$$

Hence the expected number of walks after a phase of $\tau = O(t_{\text{hit}}/\kappa)$ steps is at most $c \cdot \kappa^{100}$ for some constant $1/2 \leqslant c < 1$, and an application of Markov's inequality yields the claim. □

We are ready to prove Theorem 4.4 by combining Lemma 4.5 with Lemma 4.7.

Proof of Theorem 4.4. Let $k = \kappa^{100}$ be the number of random walks. Since we seek to reduce the number of random walks to $(\Delta/d)^{100}$, we assume in the following that that $k^{1/100} = \kappa > \Delta/d$. Further, by our precondition we have $k \leqslant \log^4 n$. Notice that we can assume that κ is bounded from below by a sufficiently large constant. Otherwise, with V_0 denoting the set of starting vertices, we have $t_{\text{coal}}(V_0) = O(t_{\text{meet}} \cdot \log(|V_0|)) = O(t_{\text{hit}})$, by Proposition 3.3 and Proposition 6.8.

We will now prove that within $O(t_{\text{hit}}/\kappa)$ steps, the number of walks reduces by a constant factor. For this purpose, we define a "dense" subset of nodes as

$$D := \left\{u \in V \colon \pi(u) \geqslant 1/(n\kappa^2)\right\}.$$

Clearly, $\pi(D) \geqslant 1 - |V \setminus D| \cdot 1/(\kappa^2 n) \geqslant 1 - \frac{1}{\kappa^2}$. Further note that for every $u \in D$, $\pi(u) \leqslant \pi_{\max} \leqslant \frac{\Delta}{nd} \leqslant k^{1/100} \cdot \frac{1}{n} = \kappa \cdot \frac{1}{n}$, where the penultimate inequality holds since $k^{1/100} > \Delta/d$ by assumption.

Lemma 4.5 with $S := D$ and Lemma 4.7 imply that a constant fraction of $k = \kappa^{100}$ random walks can be be eliminated in a single phase of $O(t_{\text{hit}}/\kappa)$ steps with constant probability > 0, provided that $\kappa > \Delta/d$.

Thus, there exists a constant $c > 0$ such that the expected time required to reduce the number of walks from k' to $\max\{k'/2, (\Delta/d)^{100}\}$ is bounded by $c \cdot t_{\text{hit}}/\sqrt[100]{k'}$, by Lemma A.7. Hence, the expected time to reduce the number of walks from $k' \leqslant \log^4 n$ to $(\Delta/d)^{100}$ is upper bounded by

$$\sum_{i=0}^{\log(\log^4 n)} \frac{c \cdot t_{\text{hit}}}{\sqrt[100]{2^i}} \leqslant c \cdot t_{\text{hit}} \sum_{i=0}^{\infty} \left(\frac{1}{\sqrt[100]{2}}\right)^i = \frac{c \cdot t_{\text{hit}}}{1 - \sqrt[100]{1/2}} = O(t_{\text{hit}}).$$

□

4.4 Bounding t_{coal} in terms of t_{hit}

In this subsection we prove the following theorem relating t_{coal} to t_{hit} (recall that $t_{\text{meet}} \leqslant 4t_{\text{hit}}$ for any graph.)

Theorem 4.8. *Let $G = (V, E)$ be any graph with maximum degree Δ and average degree d. Then $t_{\text{coal}} = O(t_{\text{hit}} + t_{\text{meet}} \cdot \log(\Delta/d))$. So in particular, for any almost-regular graph, $t_{\text{coal}} = O(t_{\text{hit}})$.*

Theorem 4.8 follows almost immediately from the previous two reductions in Theorem 4.3 and Theorem 4.4.

Proof of Theorem 4.8. By Theorem 4.3, we can reduce the number of walks from n to $O(\log^3 n)$ in $O(t_{\text{hit}})$ steps with probability at least $1 - n^{-1}$. Then, using Theorem 4.4, we can reduce the number of walks from $O(\log^3 n)$ to $(\Delta/d)^{100}$ in $O(t_{\text{hit}})$ expected time. Finally, we apply Proposition 3.3 to reduce the number of walks from $(\Delta/d)^{100}$ to 1 in $O(t_{\text{meet}} \cdot \log(\Delta/d))$ expected time to obtain the result. □

4.5 Proof of Theorem 1.3

Part (i) is was just shown in Theorem 4.8. Part (ii) follows from Part (i) together with Proposition 6.8. Finally, Part (iii) follows from the results presented in Section 7.

4.6 Conjecture and a Possible Improvement for Non-Regular Graphs

Before concluding this section, we mention an intriguing conjecture that might be useful to improve our bound on t_{coal} when $\Delta \gg d$.

Conjecture 4.9. *There exists a universal constant $C > 0$ so that for any graph $G = (V, E)$, any deterministic path $(x_0, x_1, \ldots, x_{t_{\text{hit}}})$ and any vertex $u \in V$,*

$$\sum_{t=0}^{t_{\text{hit}}} p^t_{u,x_t} \leqslant C \cdot \sum_{t=0}^{t_{\text{hit}}} \pi(x_t).$$

Note that the inequality is a stronger version than the one given in the first statement of Lemma 4.1 or Lemma 4.2. We do not know whether the conjecture is actually true in this generality. However, if it is true, it would imply that any random walk of length $4t_{\text{hit}}$ starting from an arbitrary vertex intersects with any deterministic path of length $4t_{\text{hit}}$ with constant probability > 0. This would then result in a simple proof that $t_{\text{coal}} = O(t_{\text{hit}} \cdot \log^* n)$ for any graph, since each phase of $O(t_{\text{hit}})$ steps would reduce the number of walks from k to $O(\log k)$.

Conjecture 4.9 can be also seen as the optimization problem of "predicting" a random walk $(Y_t)_{t \geqslant 0}$ for t_{hit} time steps. More precisely, we are given the start vertex of the random walk $Y_0 = u$ and for each time step $1 \leqslant t \leqslant t_{\text{hit}}$, we have to specify a vertex x_t that acts as a predictor of the random location of the random walk at step t. The goal is to maximize the (expected) number of correct predictions, which is equal to

$$\sum_{t=0}^{t_{\text{hit}}} p^t_{u,x_t}.$$

The conjecture states that regardless which prediction, i.e., which path $(x_0, x_1, \ldots, x_{t_{\text{hit}}})$ is picked, the expected number of correct predictions cannot be made larger than in the setting where the random walks starts from stationarity.

Graph	t_{meet}		t_{coal}	
General Graphs	$\Omega(1)$, $O(n^3)$	*Thm. 6.1*	$\Omega(\log n)$, $O(n^3)$	*Lem. 5.2* & *Thm. 1.3*
Regular Graphs	$\Omega(n)$, $O(n^2)$	*Thm. 6.1* & *Thm. 1.3*	$\Omega(n)$, $O(n^2)$	*Thm. 6.1* & *Thm. 1.3*
Vertex-Trans. Graphs	$\Omega(n)$, $O(n^2)$	*Thm. 6.1* & *Thm. 1.3*	$\Omega(n)$, $O(n^2)$	*Thm. 6.1* & *Thm. 1.3*

Table 2: A summary of bounds on the meeting and coalescence times graph classes. All bounds are easily shown to be tight: For general graphs the meeting time and coalescence bounds are matched by the star and the barbell graph. For vertex-transitive graphs the bounds are matched by the clique and the cycle and thus also apply to regular graphs.

One specific strategy would be to choose $x_0 = x_1 = \cdots = x_{t_{\text{hit}}} = v$ for some vertex v. In that case we know by Lemma 4.1 that

$$\sum_{t=0}^{t_{\text{hit}}} p^t_{u,v} \leqslant 4 \cdot t_{\text{hit}} \cdot \pi(v),$$

so the conjecture holds in this case.

It is also worth mentioning that we cannot replace t_{hit} by a smaller value, say, t_{mix}. Indeed if G is a two-dimensional grid, then $t_{\text{mix}} = \Theta(n)$ and choosing $x_0 = x_1 = \cdots = x_{t_{\text{hit}}} = u$, we obtain $\sum_{t=0}^{t_{\text{mix}}} p^t_{u,u} = \Omega(\log n)$, while, $\sum_{t=0}^{t_{\text{mix}}} \pi(x_t) = O(1)$.

Finally, there is some resemblance to the meeting-time-lemma in the continuous-time setting [29], however, one important difference is that in Conjecture 4.9, the right hand side depends on the actual path $(x_0, x_1, \ldots, x_{t_{\text{hit}}})$.

5 Bounding $t_{\text{coal}} \in [\Omega(\log n), O(n^3)]$

Given that worst-case upper and lower bounds have long been known for $t_{\text{mix}}, t_{\text{hit}}$ and t_{cov}, it is very natural to pose the same question for t_{meet} and t_{coal}. In the following we determine the correct asymptotic worst-case upper and lower bounds for t_{meet} and t_{coal} on (i) general graphs, (ii) regular graphs and (iii) vertex-transitive graphs. We refer to Table 2 for an overview.

5.1 General Upper Bound $t_{\text{coal}} = O(n^3)$

In this section we establish that $t_{\text{coal}} = O(n^3)$ on all graphs, which is matched for instance by the Barbell graph.

Theorem 5.1. *For any graph G we have $t_{\text{coal}} = O(n|E| \cdot \log(|E|/n))$, so in particular, $t_{\text{coal}} = O(n^3)$.*

Proof of Theorem 5.1. This follows immediately from Theorem 4.8 and the well-known bound $t_{\text{hit}} \leqslant n \cdot |E|$:

$$\begin{aligned} t_{\text{coal}} &= O(t_{\text{hit}} + t_{\text{meet}} \cdot \log(\Delta/d)) \\ &= O(t_{\text{hit}} + t_{\text{hit}} \log(\Delta/d) \\ &= O(n \cdot |E| + n \cdot |E| \cdot \log(|E|/n)) = O(n^3), \end{aligned}$$

where the last inequality holds since $|E| \leqslant n^2$. ∎

5.2 General Lower Bound $t_{\text{coal}} = \Omega(\log n)$

In this section, we prove that on any graph the coalescing time is $\Omega(\log n)$. We'll make use of the full power of vertex-dependent selector functions defined in Section 3.1. We consider a process where there is one random walk starting at each node in the graph and this random walk has the node name (number) as its identifier. The selector function we use has *home-advantage*, when one out of several random walks at a vertex must be selected to survive, the walk with the id corresponding to that node is chosen if it exists, otherwise an arbitrary one may be chosen. It is straightforward to show that the set of nodes which have an active random walk according to this process can be coupled with the coalescence process defined in Section 2.

The main task is to show that after $c \log n$ steps, for a sufficiently small c, there are at least two surviving walks in this process. In order to do this, we simply argue that there must be at least two walks that have not left their starting position. Note that there is no way for these walks to be killed, because even if other walks visited one of their 'home' nodes, the home advantage selection rule implies that the walk that began its journey at this node is the one that survived. The formal proof follows.

Lemma 5.2. *For any graph $G = (V, E)$, $|V| = n$ we have $t_{\text{coal}} = \Omega(\log n)$.*

Proof. At time 0, $S_0 = V$. For all $v \in V$ and $S \subseteq 2^{\mathbb{N}}$ define

$$\sigma_v(S) = \begin{cases} \min\{i \in S\} & \text{if } v \notin S \\ v & \text{otherwise.} \end{cases}$$

Observe that this process P' can be coupled (using Observation 3.2) with the process P defined in Section 3.1 and hence it suffices to derive a lower bound on the coalescing time for P'. Let T be the coalescence time. The subtle difference in the definition of σ_v requires every random walk to move at least once to coalesce, *i. e.,* if at least two random walks have not moved during the first t time steps, then $T > t$.

Let X_u be the indicator variable which is 1 if the random walk starting at $u \in V$ is still alive after the first $\tau = \log n/2$ time-steps, *i. e.,* $u \in S_\tau$. Let Y_i be the indicator variable which is 1 if the random walk starting at $u \in V$ did not move during the first τ time-steps. The definition of selector function family σ_v ensures that the random walk starting at u can only vanish after coalescing with a random walk on a node $v \neq u$. In particular, the process ensures independence of the Y_i. Observe that $X_u = 1$ implies $Y_u = 1$. Due to the laziness of the random walk, $\mathbb{P}[X_u = 1] \geqslant \mathbb{P}[Y_u = 1] = 1/2^\tau = 1/\sqrt{n}$.

Thus, using the independence of the Y_i,

$$\mathbb{P}[T \geqslant \tau] \geqslant \mathbb{P}\left[\sum_{u \in V} X_u \geqslant 2\right] \geqslant \mathbb{P}\left[\sum_{u \in V} Y_u \geqslant 2\right] = \mathbb{P}\left[\text{Binomial}(n, 1/\sqrt{n}) \geqslant 2\right] = \Omega(1),$$

where Binomial(n, p) denotes the binomial distribution with parameters n and p. We conclude that $E[T] = \Omega(\log n)$ which yields the claim. □

5.3 Proof of Theorem 1.5

We are now ready to put all pieces together. The upper bound on general graphs follows directly from $t_{\text{meet}} \leqslant t_{\text{coal}} = O(n^3)$, by Theorem 5.1. The lower bound on the meeting time holds by definition and the lower bound on the coalescing time follows from Lemma 5.2. For the upper bound on regular graphs we have $t_{\text{meet}} \leqslant t_{\text{coal}} = O(t_{\text{hit}}) = O(n^2)$ due to Theorem 1.3, having used the standard bound $t_{\text{hit}} = O(n^2)$ for regular graphs (see [2]). The lower bound follows from $t_{\text{coal}} \geqslant t_{\text{meet}} \geqslant t^\pi_{\text{meet}} = \Omega(n)$, by Theorem 6.1.

6 Bounding t_{meet} and Applications

Although the focus of this work is on understanding the coalescence time, in order to apply our general results, we need to devise some tools to obtain lower and upper bounds on t_{meet}. In Theorem 6.1 (Section 6.1) we establish upper and lower bounds on the meeting time in terms of $\|\pi\|_2^2 = \sum_{u\in V} \pi(u)^2$. Section 6.1 contains also a couple of other upper bounds on t_{meet} and t_{hit}. A common feature of these upper bounds is a sub-linear dependence on the spectral gap $1/(1-\lambda_2)$, which we obtain by an application of short-term upper bound on the t-step transition probabilities.

In Proposition 6.8 (Section 6.2) we establish a discrete-time counterpart of [2, Proposition 14.5], albeit with worse constants, stating that the meeting time is at most of the order of the hitting time; on vertex transitive graphs these quantities are asymptotically of the same order.

6.1 Relating Meeting Time to t_{mix} and $\frac{1}{1-\lambda_2}$

We first state some basic bounds on t_{hit} and t_{meet}, which mostly follow directly from (1) and its counterpart for the hitting times (cf. Cooper, Frieze [11]). In these bounds, we will use the following notation:

$$C_{\max} := \max_{u\in V} \sum_{t=0}^{t_{\text{mix}}} \sum_{v\in V} \left(p_{u,v}^t\right)^2,$$

$$C_{\min} := \min_{u\in V} \sum_{t=0}^{t_{\text{mix}}} \sum_{v\in V} \left(p_{u,v}^t\right)^2.$$

Note that $C_{\max}$ and $C_{\min}$ provide worst-case upper respective lower bounds on the expected collisions of two random walks of length t_{mix}, starting from the same vertex u.

Similarly, we define

$$R_{\max} := \max_{u\in V} \sum_{t=0}^{t_{\text{mix}}} p_{u,u}^t.$$

Note that $R_{\max}$ is the number of expected returns of a random walk to u during t_{mix} steps. This quantity is more convenient to bound than $C_{\max}$, for instance, it can be easily bounded by $\max_{u\in V} \pi(u) \cdot t_{\text{mix}} + \frac{1}{1-\lambda_2}$ (cf. Lemma A.1, or also [14]).

Theorem 6.1. *For any graph $G = (V, E)$, the following statements hold:*

(i) For any pair of vertices $u, v \in V$,

$$t_{\text{hit}}(u, v) \leqslant \frac{5e \cdot (\sum_{t=0}^{t_{\text{mix}}} p_{v,v}^t)}{\pi(v)}.$$

In particular, if the graph G is Γ-approximative regular, then $t_{\text{hit}}(u, v) \leqslant 5e \cdot \Gamma \cdot n \cdot \sum_{t=0}^{t_{\text{mix}}} p_{v,v}^t$.

(ii) For any pair of vertices $u, v \in V$,

$$t_{\text{meet}}(u, v) \leqslant \frac{5e^2 \cdot C_{\max}}{\|\pi\|_2^2}.$$

In particular, if the graph G is Γ-approximative regular,

$$t_{\text{meet}}(u, v) \leqslant \frac{10e^2 \cdot (4 + \log \Gamma) \cdot R_{\max}}{\|\pi\|_2^2}.$$

(iii) It holds that,

$$t^{\pi}_{\mathrm{meet}} \geqslant \frac{C_{\min}}{8\|\pi\|_2^2}.$$

In particular, if the graph G is Γ-approximative regular, $t^{\pi}_{\mathrm{meet}} = \Omega(n/\Gamma)$.

Since $C_{\min} \geqslant 1$, the last statement of the lemma implies also $t^{\pi}_{\mathrm{meet}} = \Omega(\frac{1}{\|\pi\|_2^2})$. We remark that the second upper bound on $t_{\mathrm{meet}}(u,v)$ does only depend logarithmically on Γ.

Proof. We begin by proving the first part. Consider one random walk $(X_t)_{t\geqslant 0}$, starting from an arbitrary vertex. Divide the time-interval into consecutive epochs of length $t_{\mathrm{sep}}+t_{\mathrm{mix}}$, and let $Z := \sum_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+t_{\mathrm{mix}}} \mathbf{1}_{X_t=v}$ denote the number of visits. Then, by the separation time, $\mathbb{E}[\,Z\,] \geqslant t_{\mathrm{mix}} \cdot \frac{1}{e} \cdot \pi(v)$, and the counterpart of (1) yields

$$\mathbb{P}[\,Z \geqslant 1\,] \geqslant \frac{t_{\mathrm{mix}} \cdot \frac{\pi(v)}{e}}{\mathbb{E}[\,Z \mid Z \geqslant 1\,]}.$$

Clearly, $\mathbb{E}[\,Z \mid Z \geqslant 1\,] \leqslant \max_{t=0}^{t_{\mathrm{mix}}} \sum_{s=t}^{t_{\mathrm{mix}}} p^{s-t}_{v,v} \leqslant \sum_{t=0}^{t_{\mathrm{mix}}} p^{t}_{v,v}$. Hence,

$$\mathbb{P}[\,Z \geqslant 1\,] \geqslant \frac{t_{\mathrm{mix}} \cdot \frac{\pi(v)}{e}}{\sum_{t=0}^{t_{\mathrm{mix}}} p^{t}_{v,v}} =: p.$$

This means that in every epoch of length $t_{\mathrm{sep}} + t_{\mathrm{mix}} \leqslant 5t_{\mathrm{mix}}$, the random walk has a probability of at least p to visit vertex v, and this independently of any previous epoch. Therefore, the expected number of steps until v is visited is upper bounded by

$$t_{\mathrm{hit}}(u,v) \leqslant 5t_{\mathrm{mix}} \cdot \frac{1}{p} \leqslant \frac{5e \cdot \sum_{t=0}^{t_{\mathrm{mix}}} p^{t}_{v,v}}{\pi(v)}.$$

The claim for Γ-approximative regular graphs follows from the observation that $\min_{u\in V} \pi(u) \geqslant 1/(\Gamma n)$.

We continue with the second part. Consider two independent random walks, $(X_t)_{t\geqslant 0}$, $(Y_t)_{t\geqslant 0}$ of length $t_{\mathrm{sep}} + t_{\mathrm{mix}}$ with arbitrary start vertices. Let Z be the random variable counting the number of collisions between steps t_{sep} and $t_{\mathrm{sep}} + t_{\mathrm{mix}}$, *i. e.,*

$$Z := \sum_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+t_{\mathrm{mix}}} \mathbf{1}_{X_t=Y_t}.$$

By linearity of expectation,

$$\mathbb{E}[\,Z\,] = \sum_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+t_{\mathrm{mix}}} \sum_{u\in V} \mathbb{P}[\,X_t = u\,] \cdot \mathbb{P}[\,Y_t = u\,] \geqslant t_{\mathrm{mix}} \cdot \frac{1}{e^2} \cdot \|\pi\|_2^2. \tag{14}$$

Let us now consider $\mathbb{E}[\,Z \mid Z \geqslant 1\,]$ and recall that conditioning on $Z \geqslant 1$ can be regarded as jumping to the first step $\tau := \min\{t \colon t_{\mathrm{sep}} \leqslant t \leqslant t_{\mathrm{mix}}, X_\tau = Y_\tau\}$ without knowing anything about the future steps $t > \tau$ of both walks. Therefore,

$$\mathbb{E}[\,Z \mid Z \geqslant 1\,] \leqslant \max_{u\in V} \sum_{t=0}^{t_{\mathrm{mix}}} \sum_{v\in V} \left(p^{t}_{u,v}\right)^2 = C_{\max}.$$

Plugging this into (1) and using (14) we finally arrive at

$$\mathbb{P}[Z \geqslant 1] \geqslant \frac{\mathbb{E}[Z]}{\mathbb{E}[Z \mid Z \geqslant 1]} \geqslant \frac{t_{\mathrm{mix}} \cdot \frac{1}{e^2} \cdot \|\pi\|_2^2}{C_{\max}} =: p.$$

Hence,

$$t_{\mathrm{meet}}(u,v) \leqslant 5t_{\mathrm{mix}} \cdot \frac{1}{p} \leqslant \frac{5e^2 \cdot C_{\max}}{\|\pi\|_2^2}.$$

Let us derive the result for Γ-approximate regular graphs. To this end, define

$$S_i := \left\{u \in V \colon \deg(u) \in (2^{i-1}, 2^i]\right\},$$

and note that $S_1, \dots, S_{\log n}$ forms a partition of V. Since the graph is Γ-approximate regular, at most $4 + \log(\Gamma)$ of the S_i's are non-empty. Hence there exists a set S_j with

$$\sum_{j \in S_i} \pi(j)^2 \geqslant \frac{1}{4 + \log(\Gamma)} \cdot \|\pi\|_2^2.$$

Let us now redefine Z to denote the collisions on the set S_j, *i. e.,*

$$Z := \sum_{t=t_{\mathrm{sep}}}^{t_{\mathrm{sep}}+t_{\mathrm{mix}}} \mathbf{1}_{X_t = Y_t \wedge X_t \in S_j}.$$

Then,

$$\mathbb{E}[Z] \geqslant t_{\mathrm{mix}} \cdot \sum_{j \in S_i} \frac{1}{e^2} \cdot \pi(j)^2 \geqslant t_{\mathrm{mix}} \cdot \frac{1}{e^2 \cdot (4 + \log(\Gamma))} \cdot \|\pi\|_2^2.$$

Furthermore,

$$\mathbb{E}[Z \mid Z \geqslant 1] \leqslant \max_{u \in S_i} \sum_{t=0}^{t_{\mathrm{mix}}} \sum_{v \in S_i} \left(p_{u,v}^t\right)^2 \leqslant \max_{u \in S_i} \sum_{t=0}^{t_{\mathrm{mix}}} \sum_{v \in S_i} p_{u,v}^t \cdot 2p_{v,u}^t,$$

having used reversibility, *i. e.,* $p_{u,v}^t \pi(u) = p_{v,u}^t \pi(v)$ and $\pi(v)/\pi(u) \leqslant 2$ by definition of S_i. Further,

$$\mathbb{E}[Z \mid Z \geqslant 1] \leqslant 2 \cdot \max_{u \in S_i} \sum_{t=0}^{t_{\mathrm{mix}}} \sum_{v \in V} p_{u,v}^t \cdot p_{v,u}^t \leqslant 2 \cdot \max_{u \in V} \sum_{t=0}^{t_{\mathrm{mix}}} p_{u,u}^{2t} \leqslant 2 \cdot \max_{u \in V} \sum_{t=0}^{t_{\mathrm{mix}}} p_{u,u}^t = 2 \cdot R_{\max},$$

where the last inequality holds since $p_{u,u}^t$ is non-increasing by Lemma A.1. The claim is now derived in the same way as before.

Finally, for the third statement, let $(X_t)_{t \geqslant 0}$, $(Y_t)_{t \geqslant 0}$ be two random walk starting from stationarity. Let Z be the random variable counting the number of collisions between steps 1 and $2t_{\mathrm{mix}}$, *i. e.,*

$$Z := \sum_{t=0}^{2t_{\mathrm{mix}}} \mathbf{1}_{X_t = Y_t}.$$

Then,

$$\mathbb{E}[Z] = 2t_{\mathrm{mix}} \cdot \|\pi\|_2^2.$$

In order to upper bound $\mathbb{E}[Z \mid Z \geqslant 1]$, let us write $Z = Z_1 + Z_2$ with $Z_1 := \sum_{t=0}^{t_{\text{mix}}} \mathbf{1}_{X_t = Y_t}$ and $Z_2 := \sum_{t=t_{\text{mix}}+1}^{2t_{\text{mix}}} \mathbf{1}_{X_t = Y_t}$. By Observation A.6, $\mathbb{P}[Z_1 \geqslant 1 \mid Z \geqslant 1] \geqslant \frac{1}{2}$. Therefore, by law of total expectation,

$$\mathbb{E}[Z \mid Z \geqslant 1] \geqslant \frac{1}{2} \cdot \mathbb{E}[Z \mid Z_1 \geqslant 1] \geqslant \frac{1}{2} \cdot \min_{u \in V} \sum_{t=0}^{t_{\text{mix}}} \sum_{v \in V} \left(p_{u,v}^t\right)^2 = \frac{1}{2} \cdot C_{\min}.$$

Hence,

$$\mathbb{P}[Z \geqslant 1] \leqslant \frac{2t_{\text{mix}} \cdot \|\pi\|_2^2}{\frac{1}{2} \cdot C_{\min}} =: p.$$

Divide now the time interval $\left[0, \frac{C_{\min}}{4\|\pi\|_2^2}\right]$ into consecutive time-intervals of length $2t_{\text{mix}}$ each, and let B be the number of time-interval with at least one collision. Then

$$\mathbb{E}[B] \leqslant \frac{C_{\min}}{4\|\pi\|_2^2} \cdot \frac{1}{2t_{\text{mix}}} \cdot p \leqslant \frac{1}{2}.$$

Hence by Markov's inequality, $\mathbb{P}[B \geqslant 1] \leqslant \frac{1}{2}$ and thus $t_{\text{meet}}^{\pi} \geqslant \frac{1}{2} \cdot \frac{C_{\min}}{4\|\pi\|_2^2}$. The claim for Γ-approximate regular graph follows immediately since $\|\pi\|_2^2 \leqslant \max_{u \in V} \pi(u)$, which completes the proof of the theorem. □

It is interesting to compare the upper bound on t_{meet} in Theorem 6.1 with the bound $t_{\text{meet}} = O(\frac{1}{1-\lambda_2} \cdot (\frac{1}{\|\pi\|_2^2} + \log n))$ from Cooper et al. [14, Theorem 2]. Using the trivial bound $C_{\max} \leqslant t_{\text{mix}}$ and $t_{\text{mix}} = O(\frac{\log n}{1-\lambda_2})$, we obtain $t_{\text{meet}} = O(\frac{1}{1-\lambda_2} \cdot \frac{\log n}{\|\pi\|_2^2})$, which is at most a $\log n$-factor worse. However, for certain graphs like grids or tori one may have a better control on the t-step probabilities, so that $C_{\max} \ll t_{\text{mix}}$ could be established.

Proposition 6.2. *Combining the upper bound on t_{meet} in Theorem 6.1 with the bound $t_{\text{meet}} = O(\frac{1}{1-\lambda_2} \cdot (\frac{1}{\|\pi\|_2^2} + \log n))$ from Cooper et al. [14, Theorem 2] together with Theorem 1.1 we derive*

$$t_{\text{coal}} = O\left(\frac{1}{1-\lambda_2} \cdot \frac{1}{\|\pi\|_2^2} + t_{\text{mix}} (\log n)^3\right),$$

which is at least as good as the bound of [14, Theorem 1] and equally good if one uses the trivial bound $t_{\text{mix}} = O\left(\frac{\log n}{1-\lambda_2}\right)$.

Proof. First assume $t_{\text{meet}}/t_{\text{mix}} \geqslant (\log n)^2$. In this case, by Theorem 1.1,

$$t_{\text{coal}} = O(t_{\text{meet}}) = O\left(\frac{1}{1-\lambda_2} \cdot \left(\frac{1}{\|\pi\|_2^2} + \log n\right)\right)$$

follows immediately. Next assume $t_{\text{meet}}/t_{\text{mix}} \leqslant (\log n)^2$, so $t_{\text{meet}} \leqslant t_{\text{mix}} \cdot (\log n)^2$. By Proposition 3.3, $t_{\text{coal}} = O(t_{\text{meet}} \log n) = O\left(t_{\text{mix}} (\log n)^3\right)$. □

In the following, we will try to get more concrete estimates than the ones in Theorem 6.1 by expressing the number of expected returns or $C_{\max}$ through $1 - \lambda_2$ and t_{mix}. To this end, we define

$$\beta := \min\left\{\frac{\log(1/(1-\lambda_2))}{1-\lambda_2}, t_{\text{mix}}\right\}.$$

Note that since $t_{\text{mix}} = \Omega(1/(1-\lambda_2))$, $\beta \geqslant \Omega(1/(1-\lambda_2))$. Further, $\beta \leqslant t_{\text{mix}} = O(\log n/(1-\lambda_2))$. Hence β is always sandwiched between the relaxation time $1/(1-\lambda_2)$ and mixing time.

We will frequently make use of the following result, which is a straightforward generalization of a result in the textbook by Aldous and Fill [2] from regular to Γ-approximate regular graph.

Lemma 6.3 ([2, Proposition 6.16 (iii)])**.** *Let G be any Γ-approximate regular graph. Then for any $\tau \leqslant 5n^2$,*

$$\sum_{t=0}^{\tau} p_{u,u}^t \leqslant 2\,\Gamma \cdot \sqrt{5\tau}.$$

Since $p_{u,u}^t$ is non-increasing, this implies for any $1 \leqslant t \leqslant 5n^2$, $p_{u,u}^t \leqslant \frac{20\Gamma}{\sqrt{t}}$.

As mentioned earlier, $\sum_{t=0}^{t_{\text{mix}}} p_{u,u}^t \leqslant t_{\text{mix}} \cdot \pi(u) + \frac{1}{1-\lambda_2}$ are well-known bounds. The next corollary provides an improvement in many cases:

Corollary 6.4. *For any Γ-approximate regular graph $G = (V, E)$*

$$R_{\max} = \max_{u \in V} \sum_{t=0}^{t_{\text{mix}}} p_{u,u}^t = O\left(\min\left\{\Gamma\sqrt{\frac{\log(1/(1-\lambda_2))}{1-\lambda_2}}, \Gamma^{3/2}\sqrt{t_{\text{mix}}}\right\}\right) = O\left(\Gamma^{3/2} \cdot \sqrt{\beta}\right).$$

Proof. By [2, Corollary 6.9], we have $t_{\text{mix}} \leqslant t_{\text{hit}} = O(\Gamma \cdot n^2)$. Thus, using that $p_{u,u}^t$ is non-increasing (Lemma A.1) and Corollary 6.4, we derive

$$\begin{aligned} R_{\max} = \max_{u \in V} \sum_{t=0}^{t_{\text{mix}}} p_{u,u}^t &\leqslant 2\sqrt{5}\Gamma\sqrt{t_{\text{mix}}} + \sum_{t=5n^2}^{t_{\text{mix}}} p_{u,u}^{5n^2} \leqslant O(\Gamma\sqrt{t_{\text{mix}}}) + t_{\text{mix}} \cdot 20\Gamma/n \\ &= O\left(\Gamma \cdot \sqrt{t_{\text{mix}}}\left(1 + \sqrt{t_{\text{mix}}}/n\right)\right) = O\left(\Gamma^{3/2}\sqrt{t_{\text{mix}}}\right). \end{aligned}$$

In the remainder we derive a bound in terms of $\frac{\log(1/(1-\lambda_2))}{1-\lambda_2}$. We split the expected number of returns to u at time x and obtain

$$\begin{aligned} \sum_{t=0}^{t_{\text{mix}}} p_{u,u}^t &\leqslant \sum_{t=0}^{x-1} p_{u,u}^t + \sum_{t=x}^{t_{\text{mix}}} p_{u,u}^t \\ &\leqslant \left(2\Gamma \cdot \sqrt{5x} + t_{\text{mix}} \cdot 20\Gamma/n\right) + \sum_{t=x}^{t_{\text{mix}}} \left(\pi(u) + \lambda_2^t\right) \\ &\leqslant 2\Gamma \cdot \sqrt{5x} + 21\Gamma \cdot t_{\text{mix}}/n + \frac{\lambda_2^x}{1-\lambda_2}, \end{aligned}$$

where the second inequality follows from Lemma 6.3 and Lemma A.1. Next choose $x = \frac{\ln(1/(1-\lambda_2))}{1-\lambda_2}$. Since $f(y) = y^{1/(1-y)}$ is bounded from above by $1/e$ for any $y \in (0,1)$, we have

$$\sum_{t=0}^{t_{\text{mix}}} p_{u,u}^t \leqslant 10\Gamma \cdot \sqrt{\frac{\ln\left(1/(1-\lambda_2)\right)}{1-\lambda_2}} + 21\Gamma \cdot t_{\text{mix}}/n + 1.$$

We next prove that the second term in the bound above is always asymptotically upper bounded by the first one. This is established via a simple case distinction. First, if $1/(1-\lambda_2) \leqslant n^{1/2}$, then

the claim holds because of $t_{\text{mix}} = O(\log n/(1-\lambda_2))$. Secondly, if $1/(1-\lambda_2) \geqslant n^{1/2}$, then using the same bound on t_{mix} along with the fact that for regular graphs $1-\lambda_2 = \Omega(1/n^2)$ [6]

$$t_{\text{mix}} = O\left(\frac{\log n}{\sqrt{1-\lambda_2}\cdot\sqrt{1-\lambda_2}}\right) = O\left(\frac{n\cdot 2\log(n^{1/2})}{\sqrt{1-\lambda_2}}\right) = O\left(n\cdot\sqrt{\frac{\log(1/(1-\lambda_2))}{1-\lambda_2}}\right).$$

□

Theorem 6.5. *Let G be any Γ-approximate regular graph. Then*

$$t_{\text{hit}} = O\left(n\Gamma^{5/2}\cdot\sqrt{\beta}\right),$$

and consequently,

$$t_{\text{meet}} = O\left(n\Gamma^{5/2}\cdot\sqrt{\beta}\right).$$

In particular, if Γ is constant, then

$$t_{\text{meet}} \leqslant t_{\text{coal}} = O(t_{\text{hit}}) = O\left(n\cdot\sqrt{\beta}\right).$$

The upper bound on t_{meet} and t_{coal} gives $t_{\text{meet}} = O(n^2)$ for cycles and paths (since $t_{\text{mix}} = O(n^2)$), and $t_{\text{meet}} = O(n)$ on expanders (since $1/(1-\lambda_2) = O(1)$). It thus improves the bound by Cooper et al. [14, Theorem 1], which states that for any regular graphs, $t_{\text{meet}} = O(n/(1-\lambda_2))$.

Proof. By assumption, $\pi(v) \geqslant 1/(\Gamma\cdot n)$ for $v \in V$. The bound on t_{hit} follows from the definition of $\sqrt{\beta}$, Corollary 6.4 and applying them to Theorem 6.1. Since $t_{\text{meet}} \leqslant 4t_{\text{hit}}$, the same upper bound also holds for t_{meet}. □

For any $\Gamma = O(1)$-approximate regular graph, we also improve the best-known bound on the cover time t_{cov} in terms of the eigenvalue gap, which is $t_{\text{cov}} = O\left(n\log n/(1-\lambda_2)\right)$ established by Broder and Karlin in 1989 [8].

Theorem 6.6. *For any Γ-approximate regular graph,*

$$t_{\text{cov}} = O\left(\Gamma^{5/2}\cdot n\log n\cdot\sqrt{\beta}\right).$$

Proof. This follows immediately from Theorem 6.5 and the well-known trivial bound $t_{\text{cov}} = O(t_{\text{hit}}\cdot\log n)$. □

We now derive an extension of Theorem 6.5 that is more suited for graphs with a very high degree discrepancy.

Theorem 6.7. *Let $G = (V, E)$ be any Γ-approximate regular graph. Then,*

$$t_{\text{meet}} = O\left(\frac{\Gamma^{3/2}\log(\Gamma)\cdot\sqrt{\beta}}{\|\pi\|_2^2}\right).$$

Furthermore,

$$t_{\text{meet}} = O\left(\frac{\frac{\log(\Gamma)}{1-\lambda_2} + t_{\text{mix}}\cdot\max_{u\in V}\pi(u)}{\|\pi\|_2^2}\right).$$

[6] This fact can be derived as follows. First, from [9, Lemma 1.9], we have $1-\lambda_2 \geqslant \frac{1}{\text{diam}(G)\cdot n\cdot d}$, where $\text{diam}(G)$ is the diameter of G. Further, it is shown by Feige et al. [20] (or Aldous [2, Lemma 6.10]) that $\text{diam}(G) \leqslant 3n/d$.

We point out that for constant Γ, the first statement of the theorem recovers the statement of Theorem 6.5.

Proof. As in the proof of Theorem 6.1, define $S_i := \{u \in V : \deg(u) \in (2^{i-1}, 2^i]\}$, and note that $S_1, \dots, S_{\log n}$ forms a partition of V. Since the graph is Γ-approximate regular, at most $4 + \log(\Gamma)$ of the S_i's are non-empty. Hence there exists a set S_j with

$$\sum_{j \in S_i} \pi(j)^2 \geqslant \frac{1}{4 + \log(\Gamma)} \cdot \|\pi\|_2^2.$$

We will only count collisions on that bucket, *i. e.*, $Z := \sum_{t=t_{\text{sep}}}^{t_{\text{sep}}+t_{\text{mix}}} \mathbf{1}_{X_t = Y_t} \cdot \mathbf{1}_{X_t \in S_i}$. Hence for two walks $(X_t)_t$ and $(Y_t)_t$ of length t_{mix} starting from arbitrary vertices,

$$\mathbb{E}[Z] \geqslant \frac{1}{e^2 \cdot (4 + \log(\Gamma))} \cdot t_{\text{mix}} \cdot \|\pi\|_2^2.$$

Furthermore,

$$\begin{aligned}\mathbb{E}[Z \mid Z \geqslant 1] \leqslant \max_{u \in S_i} \sum_{t=0}^{t_{\text{mix}}} \sum_{v \in S_i} \left(p_{u,v}^t\right)^2 &\leqslant 4 \max_{u \in S_i} \sum_{t=0}^{t_{\text{mix}}} p_{u,u}^{2t} \\ &= O\left(\Gamma^{3/2} \cdot \sqrt{\beta}\right)\end{aligned}$$

where the second inequality holds due to the fact that vertices in S_i have the same degree up to a factor of 2 and the final inequality holds is due to Corollary 6.4.

Plugging the two bounds into (1) yields

$$\mathbb{P}[Z \geqslant 1] = \Omega\left(\frac{t_{\text{mix}} \cdot \|\pi\|_2^2}{\Gamma^{3/2} \log(\Gamma) \cdot \sqrt{\beta}}\right) := p.$$

Hence by repeating above argument on consecutive time-intervals of length $t_{\text{sep}} + t_{\text{mix}} \leqslant 5t_{\text{mix}}$,

$$t_{\text{meet}} \leqslant \frac{1}{p} \cdot 5\, t_{\text{mix}} = O\left(\frac{\Gamma^{3/2} \log(\Gamma) \cdot \sqrt{\beta}}{\|\pi\|_2^2}\right).$$

For the second statement, we also have, by Lemma A.1,

$$4 \max_{u \in S_i} \sum_{t=0}^{t_{\text{mix}}} p_{u,u}^t \leqslant 4 \cdot \frac{1}{1 - \lambda_2} + t_{\text{mix}} \cdot \max_{u \in V} \pi(u).$$

Thus the bound on t_{meet} is derived in an analogous way.

□

6.2 Relating Meeting Time to t_{hit}

In this section we prove the following proposition which can be seen as an analogous version of [2, Proposition 14.5] in discrete time.

Proposition 6.8. *For any graph $G = (V, E)$ and $u, v \in V$ we have*

$$(\min_{u' \in V} t_{\text{hit}}(\pi, u') + t_{\text{hit}}(u, v) - t_{\text{hit}}(\pi, v))/2 \leqslant t_{\text{meet}}(u, v) \leqslant 2(\max_{u' \in V} t_{\text{hit}}(\pi, u') + t_{\text{hit}}(u, v) - t_{\text{hit}}(\pi, v)).$$

Consequently, for any graph we have $t_{\text{meet}} \leqslant 4t_{\text{hit}}$ and for any vertex transitive graph G we have $t_{\text{hit}}/2 \leqslant t_{\text{meet}} \leqslant 2t_{\text{hit}}$.

Proof. We define a pair of chains $((X_t)_{t\geqslant 0},(Y_t)_{t\geqslant 0})$ with *arbitrary* starting vertices $X_0, Y_0 \in V$, called *sequential* random walks, by

$$X_{t+1} = \begin{cases} v \in N(X_t) \text{ w.p. } \frac{1}{2|N(X_t)|} & \text{if } t \text{ is even} \\ X_t & \text{otherwise} \end{cases},$$

and

$$Y_{t+1} = \begin{cases} v \in N(Y_t) \text{ w.p. } \frac{1}{2|N(Y_t)|} & \text{if } t \text{ is odd} \\ Y_t & \text{otherwise} \end{cases}.$$

In particular, for even t (odd t, respectively) the random-walk is lazy meaning $X_{t+1} = X_t$ (and $Y_{t+1} = Y_t$, respectively).

Consider two "non-sequential" random walks $(X'_t)_{t\geqslant 0}$ and $(Y'_t)_{t\geqslant 0}$ with $X'_0 = X_0$ and $Y'_0 = Y_0$. We will couple their decisions with the walks $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$, by setting $X'_t = X_{2t}$ and $Y'_t = Y_{2t}$. Due to this coupling and since each random walk is lazy w.p. $1/2$,

$$\mathbb{P}\left[\, X'_{t+1} = Y'_{t+1} \mid X_{2t} = Y_{2t} \,\right] = \mathbb{P}\left[\, X_{2t+2} = Y_{2t+2} \mid X_{2t} = Y_{2t} \,\right] \geqslant 1/4,$$

and

$$\mathbb{P}\left[\, X'_{t+1} = Y'_{t+1} \mid X_{2t+1} = Y_{2t+1} \,\right] = \mathbb{P}\left[\, X_{2t+2} = Y_{2t+2} \mid X_{2t+1} = Y_{2t+1} \,\right] = 1/2$$

Let $t^{seq}_{\text{meet}}(u,v)$ be the meeting time of the sequential chains X_t and Y_t, *i. e.,*

$$t^{seq}_{\text{meet}}(u,v) = \min\{t \geqslant 0 \mid X_t = Y_t, X_0 = u, Y_0 = v\}$$

and $t^{seq}_{\text{meet}} = \max_{u,v} t^{seq}_{\text{meet}}(u,v)$. We seek to relate t_{meet} with t^{seq}_{meet}. Clearly, $t^{seq}_{\text{meet}}/2 \leqslant t_{\text{meet}}$ since a meeting of $X'_t = Y'_t$ implies that $X_{2t} = Y_{2t}$. For an upper bound on t_{meet} recall that X'_t and Y'_t meet, *i. e.,* $X'_t = Y'_t$ w.p. at least $1/4$ whenever $X_{2t-2} = Y_{2t-2}$ or $X_{2t-1} = Y_{2t-1}$. Hence, by independence $t_{\text{meet}} = \max_{u,v} t_{\text{meet}}(u,v) \leqslant 4(t^{seq}_{\text{meet}}/2) = 2t^{seq}_{\text{meet}}$. We conclude,

$$t^{seq}_{\text{meet}}(u,v)/2 \leqslant t_{\text{meet}}(u,v) \leqslant 2t^{seq}_{\text{meet}}(u,v). \tag{15}$$

We proceed by deriving upper and lower bounds on t^{seq}_{meet}, which gives us bounds on t_{meet}. We will make use of the following statement that is a weaker version of the original statement Aldous and Fill [2, Proposition 3.3]. For all $u, v \in V$ we have

$$\min_{u'\in V} t_{\text{hit}}(\pi,u') \leqslant t^{seq}_{\text{meet}}(u,v) - (t_{\text{hit}}(u,v) - t_{\text{hit}}(\pi,v)) \leqslant \max_{u'\in V} t_{\text{hit}}(\pi,u')$$

Using (15) we derive,

$$t_{\text{meet}}(u,v) \leqslant 2t^{seq}_{\text{meet}}(u,v) \leqslant 2(t_{\text{hit}}(u,v) - t_{\text{hit}}(\pi,v) + \max_{u'\in V} t_{\text{hit}}(\pi,u')),$$

and

$$t_{\text{meet}}(u,v) \geqslant t^{seq}_{\text{meet}}(u,v)/2 \geqslant (\min_{u'\in V} t_{\text{hit}}(\pi,u') + (t_{\text{hit}}(u,v) - t_{\text{hit}}(\pi,v)))/2,$$

which yields the first part of the claim.

For vertex transitive chains we get using $t_{\text{hit}}(\pi,u) = t_{\text{hit}}(\pi,u')$ for all $u, u' \in V$

$$t^{seq}_{\text{meet}}(u,v) = t_{\text{hit}}(u,v).$$

Thus, putting everything together and fixing $u, v \in V$ to be the nodes maximizing $t_{\text{hit}}(u,v)$, we derive

$$t_{\text{hit}} = t_{\text{hit}}(u,v) = t^{seq}_{\text{meet}}(u,v) \leqslant t^{seq}_{\text{meet}} \leqslant 2t_{\text{meet}}.$$

Similarly,

$$t_{\text{meet}} \leqslant 2t^{seq}_{\text{meet}} = 2\max_{u,v} t^{seq}_{\text{meet}}(u,v) \leqslant 2t_{\text{hit}}.$$

This yields Proposition 6.8. □

6.3 Proof of Theorem 1.4

The proof follows directly from Theorem 6.5 and Theorem 6.6.

7 Applications to Special Graphs

Here we derive t_{meet} and t_{coal} on specific topologies. Note that more general bounds for certain graph classes like regular graphs or vertex-transitive graphs have been stated earlier, see, e.g., Theorem 1.3 or Section 6.

7.1 2-Dimensional Grids/Tori and Paths/Cycles

Next we apply our machinery to the 2-dimensional grid and the 2-dimensional torus. For the continuous case a manual approach to bound meeting and coalescence times can be found in [16]. Thanks to our general results, we can not only easily derive the correct bound on t_{meet}, but also on t_{coal}. First, we recall the following well-known fact that for 2-dim. grid and torus: For any integer $t \geqslant 1$,

$$p^t_{u,u} = \pi(u) + \Omega(t^{-1}), \tag{16}$$

which can be derived, e.g., by using the central limit theorem. Further, $t_{\text{mix}} = \Theta(n)$, and combining these two results, we immediately obtain

$$C_{\min} := \min_{u\in V} \sum_{t=0}^{t_{\text{mix}}} \sum_{v\in V} (p^t_{u,v})^2 \geqslant \frac{1}{2} \min_{u\in V} \sum_{t=0}^{t_{\text{mix}}} p^{2t}_{u,u} = \Omega(\log n).$$

Thus, by Theorem 6.1.(iii), $t_{\text{meet}} = \Omega(n \log n)$. For the upper bound, we apply Theorem 1.3 together with the well-known bound $t_{\text{hit}} = O(n \log n)$ to derive $t_{\text{coal}} = O(t_{\text{hit}}) = O(n \log n)$.

For cycles or paths, the corresponding formula to (16) is

$$p^t_{u,u} = \pi(u) + \Omega(t^{-1/2}).$$

Hence $C_{\min} = \Omega(n)$, and therefore the third statement of Theorem 6.1 implies $t_{\text{meet}} = \Omega(n^2)$. For the upper bound, we apply Theorem 1.3 together with $t_{\text{hit}} = O(n^2)$ to derive $t_{\text{coal}} = O(t_{\text{hit}}) = O(n^2)$. Alternatively, the upper bound on t_{coal} could be also shown by using $t_{\text{mix}} = O(n^2)$ and applying the third statement Theorem 6.5.

7.2 d-Dimensional Grids and Tori, $d \geqslant 3$

Here the bounds on t_{meet} and t_{coal} follow immediately from our general results. First, for any regular graphs we have $t_{\text{meet}} = \Omega(n)$ (Theorem 6.1.iii). Further, it is well-known that $t_{\text{hit}} = O(n)$ (e.g., [23]), and the result follows by $t_{\text{coal}} = O(t_{\text{hit}})$ shown in Theorem 1.3. Alternatively, we could also use $t_{\text{meet}} = O(t_{\text{hit}})$ (Proposition 6.8) to deduce $t_{\text{meet}} = O(n)$. Combining this with the fact that $t_{\text{mix}} = O(n^{2/d})$ [2], Theorem 1.1 yields the correct bound $t_{\text{coal}} = O(n)$.

7.3 Hypercubes

Tight bounds for the hypercube can be obtained through different tools we provide.

Firstly, it follows trivially from Theorem 1.3: Since the hypercube is regular (in fact, it is even vertex-transitive) it suffices to consider the hitting time. We have $t_{\text{hit}} = O(n)$ (see e.g. [25]) and recall that $t_{\text{hit}} = \Omega(n)$ by Theorem 6.1. Hence applying Theorem 1.3 yields $t_{\text{coal}} = \Theta(t_{\text{hit}}) = \Theta(n)$.

Alternatively, we could also use the more elementary bound $t_{\text{meet}} \leqslant 2t_{\text{hit}}$ by Proposition 6.8 to conclude $t_{\text{meet}} = O(n)$. Since it is a well-known fact that $t_{\text{mix}} = O(\log n \cdot \log\log n)$ [23], we obtain by Theorem 1.1 that $t_{\text{coal}} = \Theta(t_{\text{meet}}) = \Theta(n)$.

7.4 (Regular) Expanders

It is not surprising that on regular expander graphs, we have $t_{\text{coal}} = \Theta(n)$ and there is a multitude of approaches to establish this (for instance, the result is a consequence of the main result by [14]). With regard to our bounds, the easiest route is to follow the arguments for the hypercube: Combine the result $t_{\text{hit}} = O(n)$ (e.g., [8]) together with our bound $t_{\text{coal}} = O(t_{\text{hit}})$ (Theorem 1.3). The lower bound $t_{\text{coal}} \geqslant t_{\text{meet}} = \Omega(n)$ holds for any regular graph.

7.5 Real World Graph Models

There is a variety of different graph models for "real world" networks. In this subsection we demonstrate that random walks coalesce quickly on these graphs by establishing several bounds on t_{coal} which are sublinear in n.

First note that common features of real world graph models are (i) a power law degree distribution with exponent $\beta \in (2,3)$ and (ii) high expansion, *i. e.,* $1 - \lambda_2$ is not too large, and hence $t_{\text{mix}} = O(\log n)$. Notice that (i) $\beta \in (2,3)$ implies that w.h.p. we have $\Delta = O(n^{1-\varepsilon})$, and hence $\|\pi\|_2^2 \leqslant \max_{u \in V} \pi(u) \leqslant n^{-\varepsilon}$, for $\varepsilon > 0$.

For the sake of concreteness, let us take a specific model by Gkantsidis, Mihail and Saberi [21], which was also analyzed by Cooper et al [14]. In this model, for some $\alpha \in (2,3)$ we generate a random graph which has $\Theta(n/d^\alpha)$ vertices of degree d and an eigenvalue gap $1/(1-\lambda_2) = O(\log^2 n)$. Cooper et al. [14] derived the general bound $t_{\text{coal}} = O(\frac{1}{1-\lambda_2} \cdot (\|\pi\|_2^2 + \log^4 n))$, which implies $t_{\text{coal}} = O(n^{(\alpha-1)/2} \cdot \log^2 n)$ - a sublinear bound on the coalescing time. However, this leaves open how close t_{coal} and t_{meet} are.

Combining Theorem 1.1 with the fact that $t_{\text{mix}} = O(\log n/(1-\lambda_2)) = O(\log^3 n)$, we immediately obtain $t_{\text{coal}} = \Theta(t_{\text{meet}})$, without having to know the actual value of t_{meet} [7].

More generally, we have the following result, saying that we have $t_{\text{coal}} = \Theta(t_{\text{meet}})$ whenever t_{mix} is slightly smaller than $1/\|\pi\|_2^2$:

Theorem 7.1. *Let $G = (V, E)$ be any graph. Then,*

$$t_{\text{coal}} = O\left(t_{\text{meet}} \cdot \left(1 + \sqrt{t_{\text{mix}} \cdot \|\pi\|_2^2} \cdot \log n\right)\right).$$

In particular, whenever $t_{\text{mix}} \cdot \log^2 n \leqslant 1/\|\pi\|_2^2$, we have $t_{\text{coal}} = \Theta(t_{\text{meet}})$.

Proof. First, by the third statement of Theorem 6.1, we have $t_{\text{meet}} \geqslant \frac{1}{9\|\pi\|_2^2}$. Inserting this into Theorem 1.1 yields the upper bound. The lower bound for the setting $t_{\text{mix}} \cdot \log^2 n \leqslant 1/\|\pi\|_2^2$ trivially holds since $t_{\text{coal}} \geqslant t_{\text{meet}}$. □

[7] That being said, deriving the correct bound on t_{meet} is an interesting open problem. So far, it seems rather difficult to use one of our "off-the-shelf" bounds or the results from [14]. One potential route towards a tight bound may involve stronger bounds on $R_{\max}$, as suggested by the second upper bound on $t_{\text{meet}}(u, v)$ in Theorem 6.1.

It is worth comparing this result with the bound derived by Cooper et al. [14]:

$$t_{\text{coal}} = O\left(\frac{1}{1-\lambda_2} \cdot \left(\frac{1}{\|\pi\|_2^2} + \log^4 n\right)\right). \tag{17}$$

The advantage of (17) is that requires relatively little knowledge about G; only $\frac{1}{1-\lambda_2}$ and $\frac{1}{\|\pi\|_2^2}$ (which is equivalent to knowing the degree distribution) are needed. One potential drawback of the bound in (17) however, is that it involves the product of two factors $\frac{1}{1-\lambda_2}$ and $\frac{1}{\|\pi\|_2^2}$, each of which is a lower bound on the meeting time on its own. For instance for regular graphs, by Theorem 6.5, we immediately obtain that $t_{\text{coal}} = O(t_{\text{meet}} \log n) = O\left(n \log n \cdot \sqrt{\frac{1}{1-\lambda_2} \cdot \log(\frac{1}{1-\lambda_2})}\right)$. As a consequence, for regular graphs, our bound improves over the bound in (17) whenever $\frac{1}{1-\lambda_2} \geqslant \log^{2+\varepsilon} n$ for an arbitrarily small constant $\varepsilon > 0$.

It is also interesting to consider an alternative graph model for real world networks, proposed by Mihail et al. [27]. Also in this model, the degree distribution has the same Power law with exponent $\alpha \in (2, 3)$, but there is a stronger bound on the spectral gap, $\frac{1}{1-\lambda_2} = O(1)$ [27]. Hence Theorem 7.1 implies $t_{\text{coal}} = \Theta(t_{\text{meet}})$. Further, thanks to Proposition 6.2 (or alternatively, the bound by Cooper et al. (17)) we get the explicit bound $t_{\text{coal}} = O(1/\|\pi\|_2^2)$, which is asymptotically tight due to the trivial lower bound $t_{\text{meet}} = \Omega(1/\|\pi\|_2^2)$ (Theorem 6.1).

7.6 Binary Trees

In this subsection, we derive a lower bound $t_{\text{meet}} = \Omega(n \log n)$ for complete binary trees. Unfortunately, this bound does not follow directly from our general results and a manual analysis is required. To some extent, this is due to the structural difference between nodes close to the leaves and nodes close to the root. While a collision close to the leaves triggers $\Theta(n \log n)$ expected additional collisions, a collision near the root triggers only $\Theta(n)$ additional collisions.

Our proof consists of the following two steps. In Section 7.6.1, we first provide a lower bound on the probability that a random walk starting from any nodes $u \in V$ is on a leaf after $O(\log n)$ steps. We also show that any t-step probability $p_{u,v}^t$ is at the most return probability for a leaf. Both results shown in Lemma 7.3 are derived by projecting the random walk on the tree to a random walk on a weighted path of length $\log_2 n - 1$.

In Section 7.6.2, we proceed to analyzing the expected number of collisions between two random walks in n steps. The main component is Lemma 7.5, establishing that this number is at least $\Omega(\log n)$ provided the walks start from the same vertex not too far from the root. This result is complemented by a union-bound type argument in Lemma 7.4, showing that it is unlikely that two random walks collide on a vertex close to the root. Combining the two results and applying them to (1) establishes the desired lower bound $t_{\text{meet}} = \Omega(n \log n)$.

For the other bounds on t_{meet} and t_{coal}, we combine $t_{\text{coal}} = O(t_{\text{hit}})$ (Theorem 1.3) with the well-known fact $t_{\text{hit}} = O(n \log n)$ (cf. [2]) to obtain $t_{\text{coal}} = O(n \log n)$. Together with the established lower bound, this shows that t_{meet} and t_{coal} are both of order $\Theta(n \log n)$.

7.6.1 Bounds on the t-step probabilities

We assume that the complete binary tree has $\log_2 n - 1$ levels, *i. e.,* there are $n/2$ leaves and the total number of nodes is $n-1$. We define $\mathcal{L} \subseteq V$ to be the set of leaves. For the analysis, it will be helpful to relate a random walk on the binary tree to a corresponding random walk on a weighted path $\widetilde{G}$ of length $\log_2 n - 1$ with nodes $\widetilde{V} = \{1, 2, \ldots, \log_2 n - 1\}$, where each vertex on the path corresponds

to all vertices in the binary tree on the same level. Let Q denote the $(\log_2 n - 1) \times (\log_2 n - 1)$ transition matrix of the corresponding weighted random walk. For $i, j \in \widetilde{V}$ we have

$$q_{ij} = \begin{cases} 1/2 & \text{if } i = j \\ 1/3 & \text{if } j > i, i \neq 1 \\ 1/2 & \text{if } j > i, i = 1 \\ 1/6 & \text{if } j < i, i \neq \log_2 n - 1 \\ 1/2 & \text{if } j < i, i = \log_2 n - 1 \ . \end{cases}$$

Let $\widetilde{\pi}$ denote the stationary distribution of this process. Since the random walk on $\widetilde{G}$ is also lazy, Lemma A.1 implies that $q^t_{\log_2 n - 1, \log_2 n - 1} \geqslant \frac{1}{4}$ for all $t \geqslant 0$.

Define $\tau_{u,\mathcal{L}} := \min_{v \in \mathcal{L}} \tau_{u,v}$ to be the random variable to the first time-step random visits a leaf where the walk starts at u; $\tau_{u,\mathcal{L}} = 0$ if u is a leaf. We will frequently use the following two simple facts about random walks on binary trees:

Lemma 7.2 ([28])**.** *Let G be any tree, and u and v be two adjacent nodes. Then $t_{\text{hit}}(u, v) = 2 \cdot n_{u,v} - 1$, where $n_{u,v}$ is the number of vertices in the subtree containing u obtained by deleting the edge $\{u, v\}$.*

Lemma 7.3. *Let G be a complete binary tree, and let u be an arbitrary node. Then the following statements hold: there is a constant $c_1 > 0$, so that*

1. *for any $t \geqslant c_1 \log_2 n$, $p^t_{u,\mathcal{L}} \geqslant \frac{1}{5}$. Moreover, if u is a leaf, then the same inequality holds for all $t \geqslant 0$.*
2. *For any vertex $v \in V$, $p^t_{u,v} \leqslant 7 \cdot p^{t - 2c_1 \log_2 n}_{w,w}$, where w is any leaf.*

Proof. To derive a lower bound on $p^t_{u,\mathcal{L}}$, we consider the contracted binary tree $\widetilde{G}$. Recall that $\tau_{u,\mathcal{L}}$ is the random variable of the first time-step at which the random walk visits a leaf where the random walk starts at u; $\tau_{u,\mathcal{L}} = 0$ if u is a leaf. By conditioning on the first visit to a leaf,

$$p^t_{u,\mathcal{L}} \geqslant \sum_{s=0}^{2c_1 \log_2 n} \mathbb{P}\left[\, \tau_{u,\mathcal{L}} = s \,\right] \cdot q^{t-s}_{\mathcal{L},\mathcal{L}} \geqslant \frac{1}{4} \cdot \left(1 - \mathbb{P}\left[\, \tau_{u,\mathcal{L}} > 2c_1 \log_2 n \,\right]\right),$$

since the chain Q can be seen as a projection of P to the line. Our next claim is that $\mathbb{P}\left[\, \tau_{u,\mathcal{L}} > c_1 \log_2 n \,\right] \leqslant n^{-2}$, provided that the constant $c_1 > 0$ is sufficiently large. This can be derived by coupling the random walk on a binary tree, starting from the root, with a random walk on the integers, starting from zero and waiting until the random walk reaches the vertex $\log_2 n - 1$. Using a Chernoff bound for $X := \sum_{i=1}^{c_1 \log_2 n} X_i$, with $\mathbb{P}\left[\, X_i = +1 \,\right] = 1/3$, $\mathbb{P}\left[\, X_i = -1 \,\right] = 1/6$ and $\mathbb{P}\left[\, X_i = 0 \,\right] = 1/2$, we conclude that $\mathbb{P}\left[\, \tau_{u,\mathcal{L}} > c_1 \log_2 n \,\right] \leqslant n^{-2}$, which implies the first statement.

To prove the second statement, consider the first $c_1 \log_2 n$ steps of a random walk starting at u. Similarly as before,

$$p^t_{u,v} \leqslant \sum_{s=0}^{c_1 \log_2 n} \mathbb{P}\left[\, \tau_{u,\mathcal{L}} = s \,\right] \cdot \max_{w \in \mathcal{L}} p^{t-s}_{w,v} + \mathbb{P}\left[\, \tau_{u,\mathcal{L}} > 4 \log_2 n \,\right].$$

As seen above, $\mathbb{P}\left[\, \tau_{u,\mathcal{L}} > c_1 \log_2 n \,\right] \leqslant n^{-2}$ and therefore

$$p^t_{u,v} \leqslant \max_{t - c_1 \log_2 n \leqslant s \leqslant t} \ \max_{w \in \mathcal{L}} p^s_{w,v} + n^{-2}. \tag{18}$$

Let us now compare $p^t_{w,v}$ to $p^t_{v,w}$. Since the random walk is time-reversible, we have

$$p^t_{w,v} \cdot \pi(w) = p^t_{v,w} \cdot \pi(v),$$

and hence

$$p^t_{u,v} \leqslant \max_{t - c_1 \log_2 n \leqslant s \leqslant t} \max_{w \in \mathcal{L}} 3 \cdot p^s_{v,w} + n^{-2}.$$

Applying (18) to each $p^s_{v,w}$, we conclude that

$$p^t_{u,v} \leqslant 6 \cdot \max_{t - 2c_1 \log_2 n \leqslant s \leqslant t} \max_{w,w' \in \mathcal{L}} p^s_{w,w'} + 4n^{-2}.$$

Further, by symmetry $p^s_{w,w'}$ is maximized if $w = w'$, so that

$$p^t_{u,v} \leqslant 6 \cdot \max_{t - 2c_1 \log_2 n \leqslant s \leqslant t} p^s_{w,w} + 4n^{-2},$$

where w is any leaf. Applying Lemma A.1, it follows that the maximum is attained for $s = t - 2c_1 \log_2 n$ and $p^s_{w,w} \geqslant \frac{1}{2n-2}$, which implies the second statement. □

7.6.2 Establishing the Lower Bound on the Meeting Time

We now prove that the meeting time on binary trees is $\Omega(n \log n)$. The intuition for this is as follows. While two random walks of length $\Theta(n \log n)$ will lead to $\Theta(\log n)$ expected collisions, it turns out that the distribution of collisions is poorly concentrated. In fact we will prove that, conditional on the existence of at least one collision, the expected number of total collisions is $\Omega(\log n)$. This will imply the desired lower bound on the meeting time. A slight complication is that the collision could occur on different nodes, which is why we will first bound the probability for a collision to occur close to the root.

Let us define U to be the set of all nodes that have distance at least $\frac{1}{2} \log_2 n$ from the root. Note that $|V \setminus U| \leqslant 2 \cdot \sqrt{n}$. Further, let $\mathcal{E}$ denote the event that two random walks starting from the stationary distribution of length $n \log_2 n$ meet on a vertex in $V \setminus U$.

Lemma 7.4. *We have* $\mathbb{P}[\mathcal{E}] \leqslant n^{-1/3}$.

Proof. By the Union Bound, $\mathbb{P}[\mathcal{E}] \leqslant \sum_{t=1}^{n \log_2 n} \sum_{u \in V \setminus U} \pi(u)^2 = n \log_2 n \cdot 2\sqrt{n} \cdot \left(\frac{2}{n}\right)^2 \leqslant \frac{8 \log_2 n}{\sqrt{n}}$. □

Lemma 7.5. *For any node* $u \in U$ *and any* $2\sqrt{n} \leqslant t \leqslant n/200$, *we have* $\sum_{v \in \mathcal{L}} (p^t_{u,v})^2 = \Omega(1/t)$.

Proof. Recall that $\sum_{v \in \mathcal{L}} (p^t_{u,v})^2$ is the probability of two non-interacting, independent random walks starting from u to meet at the same leaf at time t.

Our first claim is that with probability at least $c_1 > 0$, both random walks reach a leaf before returning to u within $4 \log_2 n$ steps. To prove this claim, recall that with probability at least $1 - 2n^{-2}$, both random walks reach a leaf before step $4 \log_2 n$. Secondly, by [25, Proposition 2.3], applied to the collapsed binary tree $\widetilde{G}$, where node u is at level ℓ and the leafs are in level $\log_2 n - 1$, it follows that the probability that a random walk starting at u visits a leaf before returning to u is

$$p := \frac{1}{(t^{\widetilde{G}}_{\text{hit}}(\ell, \log_2 n - 1) + t^{\widetilde{G}}_{\text{hit}}(\log_2 n - 1, \ell)) \cdot \pi^{\widetilde{G}}_\ell}.$$

Further, $t^{\widetilde{G}}_{\text{hit}}(\ell, \log_2 n - 1) = O(\log n)$ and $t^{\widetilde{G}}_{\text{hit}}(\log_2 n - 1, \ell) \leqslant \min_{w \in \mathcal{L}} t_{\text{hit}}(w, u) \leqslant 2^{\log_2 n - \ell + 2}$ by Lemma 7.2 and $\pi^{\widetilde{G}}_\ell \leqslant 2^{-\log_2 n + \ell + 2}$, where $\text{dist}(\mathcal{L}, u) = \min_{w \in \mathcal{L}} \text{dist}(w, u)$. Hence p is at least some

constant > 0. Hence with probability at least $p^2 - 2n^{-1}$, both random walks reach a leaf before time $4\log_2 n$ without returning to u.

Consider now the original binary tree, and one of the two random walks starting from a leaf w at some time $\in [1, 4\log_2 n]$ up until timestep $t - 4\log_2 n$. Consider the shortest path from w to the root, and let z be a node that is on this shortest path and has distance $\log_2(100t)$ from w. Applying Lemma 7.2, it follows that $t_{\text{hit}}(w,z) = \sum_{i=1}^{\log_2(100t)} 2^i - 1 \geqslant 49t$. By Markov's inequality,

$$\mathbb{P}\left[\,\tau(w,z) \geqslant 2t_{\text{hit}}(w,z)\,\right] \leqslant 1/2. \tag{19}$$

Now divide the random walk into consecutive epochs of length $2t_{\text{hit}}(w,z) + 4\log_2 n$. Combining (19) and Lemma 7.3 it follows that the random walk will visit the vertex z in each epoch with probability at least $\frac{1}{2}$, conditional on having not visited the vertex z in any of the previous epochs. Therefore for any integer $\lambda \geqslant 1$,

$$\mathbb{P}\left[\,\tau(w,z) \geqslant \lambda \cdot 3 \cdot t_{\text{hit}}(w,z))\,\right] \leqslant \mathbb{P}\left[\,\tau(w,z) \geqslant \lambda \cdot (2 \cdot t_{\text{hit}}(w,z) + 4\log_2 n)\,\right] \leqslant 2^{-\lambda},$$

where the first inequality holds since $t_{\text{hit}}(w,z) = \Omega(t) = \omega(\log n)$. Hence,

$$t_{\text{hit}}(w,z) \leqslant \frac{1}{20}t_{\text{hit}}(w,z) + \mathbb{P}\left[\tau(w,z) \geqslant \frac{1}{20}t_{\text{hit}}(w,z)\right] \cdot c_1 \cdot 3t_{\text{hit}}(w,z) + \sum_{\lambda=c_1}^{\infty} 2^{-\lambda} \cdot \lambda \cdot 3t_{\text{hit}}(w,z),$$

and it follows that by choosing the constant $c_1 > 0$ large enough, there is a constant $c_2 = c_2(c_1) > 0$ so that

$$\mathbb{P}\left[\tau(w,z) \geqslant \frac{1}{10}t_{\text{hit}}(w,z)\right] \geqslant c_2.$$

Hence with probability at least $c_2 > 0$, the random walk does not reach the node z before time t. Further, with probability at least $1 - n^{-2}$, the random walk visits a leaf, say, w', before step t, say at step s, and therefore by Lemma 7.3, the random walk is at a leaf at step t with probability at least $(1 - n^{-2}) \cdot \frac{1}{5}$.

Now define $\mathcal{F}_1 := \{\tau(w,z) \geqslant t\}$ and $\mathcal{F}_2 := \{X_t \in \mathcal{L}\}$. Clearly, the events $\mathcal{F}_1$ and $\mathcal{F}_2$ are positively correlated so that

$$\mathbb{P}\left[\,\mathcal{F}_1 \cap \mathcal{F}_2\,\right] \geqslant \mathbb{P}\left[\,\mathcal{F}_1\,\right] \cdot \mathbb{P}\left[\,\mathcal{F}_2\,\right] \geqslant c_2 \cdot \left(1 - n^{-2}\right) \cdot \frac{1}{5}.$$

Combining all the events, we conclude that with constant probability $c_3 > 0$ both random walks are on a leaf at step t and have never left the subtree with root z. For one walk, the distribution will be uniform over all the leafs within a subtree whose root is the vertex closest to the root ever visited. Hence let $\mathcal{L}_1$ be all the leafs that have a non-zero probability to be visited at step t by the first random walk, and $\mathcal{L}_2$ similarly. W.l.o.g. let $|\mathcal{L}_1| \leqslant |\mathcal{L}_2|$ and observe that $\mathcal{L}_1 \subseteq \mathcal{L}_2$ since both walks start at the same node. Therefore,

$$\sum_{v \in \mathcal{L}} (p_{u,v}^t)^2 \geqslant \left(p^2 - 2n^{-1}\right) \cdot c_3 \cdot \sum_{v \in \mathcal{L}_1 \cap \mathcal{L}_2} \frac{1}{|\mathcal{L}_1|} \cdot \frac{1}{|\mathcal{L}_2|} \geqslant \left(p^2 - 2n^{-1}\right) \cdot c_3 \cdot \frac{1}{|\mathcal{L}_2|} \geqslant \left(p^2 - 2n^{-1}\right) \cdot c_3 \cdot \frac{1}{10t},$$

where the last inequality holds since we are conditioning on the event that none of the two random walks reaches the vertex z. □

Theorem 7.6. *For the binary tree it holds that* $t_{\text{meet}} = \Omega(n \log n)$.

Proof. We first only consider collisions on nodes in U by two random walks $(X_t)_{t\geqslant 0}$, $(Y_t)_{t\geqslant 0}$ starting from the stationarity distribution. More formally, we are interested in the random variable

$$Z := \sum_{t=1}^{cn\log_2 n} \sum_{v\in V\setminus U} \mathbf{1}_{X_t=Y_t=v}.$$

By linearity of expectations,

$$\mathbb{E}[Z] \leqslant \sum_{t=1}^{cn\log_2 n} \sum_{v\in V} \mathbb{P}[X_t = Y_t = v] \leqslant cn\log_2 n \cdot \sum_{v\in V} \pi(v)^2 \leqslant n\log_2 n \cdot n \cdot \left(\frac{2}{n}\right)^2 = 4c\log_2 n,$$

and clearly,

$$\mathbb{P}[Z \geqslant 1] = \frac{\mathbb{E}[Z]}{\mathbb{E}[Z \mid Z \geqslant 1]}.$$

Hence to derive an upper bound on $\mathbb{P}[Z \geqslant 1]$, we will derive a lower bound on $\mathbb{E}[Z \mid Z \geqslant 1]$. In order to this, it will be helpful if we can work under the assumption that the first collision occurs in the first half of the walk. To this end, let Z_1 be the indicator random variable that is 1 if a collision appears before time $\frac{1}{2}n\log_2 n$ and Z_2 be the indicator random variable if a collision appears after time $\frac{1}{2}n\log_2 n$. Applying Observation A.6, we get $\mathbb{P}[Z_1 \geqslant 1 \mid Z \geqslant 1] \geqslant \frac{1}{2}$. Thus,

$$\begin{aligned}
\mathbb{E}[Z \mid Z \geqslant 1] &\geqslant \mathbb{P}[Z_1 \geqslant 1 \mid Z \geqslant 1] \cdot \mathbb{E}[Z \mid Z_1 \geqslant 1] \\
&\geqslant \frac{1}{2} \cdot \sum_{t=0}^{t_{\mathrm{mix}}} \sum_{v\in V} \left(p_{u,v}^t\right)^2 \\
&\geqslant \frac{1}{2} \cdot \sum_{t=2\sqrt{n}}^{t_{\mathrm{mix}}} \sum_{v\in\mathcal{L}} (p_{u,v}^t)^2 \geqslant \sum_{t=2\sqrt{n}}^{t_{\mathrm{mix}}} \Omega(1/t) = c_4\log_2 n,
\end{aligned}$$

where $c_4 > 0$ is a constant.

Consequently, for the modified process where collisions are only allowed on nodes in U

$$\mathbb{P}[Z \geqslant 1] \leqslant \frac{4c\log_2 n}{c_4\log_2 n}.$$

Hence by choosing $c = \min\{c_4/8, 1\}$, it follows that $\mathbb{P}[Z \geqslant 1] \leqslant 1/2$.

This implies that with probability at least $1/2$, no collision occurs on nodes in U before time $cn\log_2 n$. Furthermore, by Lemma 7.4, with probability at least $1 - n^{-1/3}$ there is no collision on nodes in $V \setminus U$. Therefore, with probability at least $1/2 - n^{-1/3}$, there is no collision among two random walks before time $cn\log_2 n$, and we have shown that $t_{\mathrm{meet}} = \Omega(n\log n)$. □

7.7 Star

Clearly, the coalescing time is $\Theta(\log n)$ which could be easily shown by a direct analysis. For the sake of completeness, we point out that the upper bound also follows from Proposition 3.3 and the fact that $t_{\mathrm{meet}} = O(1)$. The matching lower bound follows from the general bound $t_{\mathrm{coal}} = \Omega(\log n)$, holding for *any graph* (see Lemma 5.2).

Acknowledgments.

The authors would like to thank Petra Berenbrink, Robert Elsässer, Nikolaos Fountoulakis, Peter Kling, and Perla Sousi for helpful discussions.

References

[1] D. Aldous. Meeting times for independent markov chains. *Stochastic Processes and their Applications*, 38(2):185–193, 1991.

[2] D. Aldous and J. Fill. Reversible markov chains and random walks on graphs, 2002. Unpublished. http://www.stat.berkeley.edu/~aldous/RWG/book.html.

[3] N. Alon, C. Avin, M. Koucký, G. Kozma, Z. Lotker, and M. R. Tuttle. Many random walks are faster than one. *Combinatorics, Probability & Computing*, 2011.

[4] M. Barlow, Y. Peres, and P. Sousi. Collisions of random walks. *Annales de l' Institut Henri Poincare*, 48(4):922–946, 2012.

[5] L. Becchetti, A. Clementi, E. Natale, F. Pasquale, and R. Silvestri. Plurality Consensus in the Gossip Model. In *SODA 15*, pages 371–390, 2015.

[6] L. Becchetti, A. E. F. Clementi, E. Natale, F. Pasquale, and L. Trevisan. Stabilizing consensus with many opinions. In *SODA*, pages 620–635, 2016.

[7] P. Berenbrink, G. Giakkoupis, A.-M. Kermarrec, and F. Mallmann-Trenn. Bounds on the voter model in dynamic networks. In *ICALP*, 2016.

[8] A. Z. Broder and A. R. Karlin. Bounds on the cover time. *Journal of Theoretical Probability*, 2:101120, 1989.

[9] F. Chung. *Spectral Graph Theory*. AMS, 1992.

[10] K.-M. Chung, H. Lam, Z. Liu, and M. Mitzenmacher. Chernoff-Hoeffding Bounds for Markov Chains: Generalized and Simplified. In *STACS*, pages 124–135, 2012.

[11] C. Cooper and A. M. Frieze. The cover time of random regular graphs. *SIAM J. Discrete Math.*, 18(4):728–740, 2005.

[12] C. Cooper and N. Rivera. The Linear Voting Model. In *ICALP*, volume 55, pages 144:1–144:12, 2016.

[13] C. Cooper, A. Frieze, and T. Radzik. Multiple random walks in random regular graphs. *SIAM J. Discret. Math.*, 23(4):1738–1761, 2009.

[14] C. Cooper, R. Elsässer, H. Ono, and T. Radzik. Coalescing Random Walks and Voting on Connected Graphs. *SIAM Journal on Discrete Mathematics*, 27(4):1748–1758, 2013.

[15] C. Cooper, R. Elsässer, and T. Radzik. The power of two choices in distributed voting. In *ICALP*, pages 435–446, 2014.

[16] J. Cox. Coalescing random walks and voter model consensus times on the torus in $\mathbb{Z}^d$. *The Annals of Probability*, 17(4):1333–1366, 1989.

[17] A. Czumaj and C. Sohler. Testing expansion in bounded-degree graphs. *Combinatorics, Probability & Computing*, 19(5-6):693–709, 2010.

[18] K. Efremenko and O. Reingold. How well do random walks parallelize? In *RANDOM*, pages 476–489, 2009.

[19] R. Elsässer and T. Sauerwald. Tight bounds for the cover time of multiple random walks. *Theor. Comput. Sci.*, 412(24):2623–2641, 2011.

[20] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized Broadcast in Networks. *Random Structures & Algorithms*, 1(4):447–460, 1990.

[21] C. Gkantsidis, M. Mihail, and A. Saberi. Conductance and congestion in power law graphs. In *SIGMETRICS*, pages 148–159, 2003.

[22] Y. Hassin and D. Peleg. Distributed probabilistic polling and applications to proportionate agreement. *Information and Computation*, 171(2):248–268, 2001.

[23] D. Levin, Y. Peres, and E. Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2006.

[24] P. Lezaud. Chernoff-type bound for finite markov chains. *Ann. Appl. Probab.*, 8(3):849–867, 1989.

[25] L. Lovász. Random walks on graphs: A survey. *Combinatorics, Paul Erdős is Eighty*, 2:1–46, 1993.

[26] A. W. Marcus, D. A. Spielman, and N. Srivastava. Interlacing families iv: Bipartite ramanujan graphs of all sizes. In *FOCS*, pages 1358–1377, 2015.

[27] M. Mihail, C. H. Papadimitriou, and A. Saberi. On certain connectivity properties of the internet topology. *J. Comput. Syst. Sci.*, 72(2):239–251, 2006.

[28] J. Moon. Random walks on random trees. *J. Austral. Math. Soc.*, 15, 1973.

[29] R. I. Oliveira. On the coalescence time of reversible random walks. *Trans. Am. Math. Soc.*, 364(4):2109–2128, 2012.

[30] R. I. Oliveira. Mean field conditions for coalescing random walks. *Ann. Probab.*, 41(5):3420–3461, 2013.

[31] Y. Peres and P. Sousi. Mixing times are hitting times of large sets. *Journal of Theoretical Probability*, 28(2):488–519, 2015.

A Basic Results about Markov Chains

We will frequently use the following basic fact about lazy random walks, which in fact also holds for arbitrary reversible Markov chains:

Lemma A.1 (cf. [23, Chapter 12])**.** *Let P be the transition matrix of a reversible Markov chain with state space Ω. Then the following statements hold:*

(i) If P is irreducible, then for any two states $x, y \in \Omega$,

$$p^t_{x,y} \leqslant \pi(y) + \sqrt{\frac{\pi(y)}{\pi(x)}} \cdot \lambda^t,$$

where $\lambda := \max\{\lambda_2, |\lambda_n|\}$ and $\lambda_1 \geqslant \lambda_2 \geqslant \cdots \geqslant \lambda_n$ are the n real eigenvalues of the matrix P.

(ii) If the Markov chain is a non-lazy random walk on a bipartite regular graph with two partitions V_1 and V_2, then for any pair of states x, y in the same partition

$$p^t_{x,y} \leqslant \frac{1}{n} \cdot \left(1 + (-1)^{t-1}\right) + \max\{\lambda_2, |\lambda_{n-1}|\}.$$

Similarly, if x and y are in opposite partitions,

$$p^t_{x,y} \leqslant \frac{1}{n} \cdot \left(1 + (-1)^{t}\right) + \max\{\lambda_2, |\lambda_{n-1}|\}.$$

(iii) If the Markov chain is lazy, then for any state $x \in \Omega$, $p^t_{x,x}$ is non-increasing in t.

Proof. The first statement can be found in [23, Equation 12.11].

For the second statement, recall the spectral representation [23, Lemma 12.2 (iii)]

$$p^t_{x,y} = \pi(y) + \pi(y) \cdot \sum_{k=2}^{n} u_k(x) \cdot u_k(y) \cdot \lambda_k^t, \tag{20}$$

where u_k is the corresponding eigenvector to λ_k. Since all eigenvalues are non-negative, we conclude from (20) that $p^t_{x,x}$ is non-increasing in t as needed. Since G is bipartite and regular, it is not difficult to verify that $\lambda_n = -1$ and $u_n(x) = \sqrt{1/n}$ if $x \in V_1$ and $u_n(x) = -\sqrt{1/n}$ if $x \in V_2$ is the corresponding eigenvector. Hence,

$$\begin{aligned}
\left| p^t_{x,y} - \frac{1}{n} \cdot \left(1 + (-1)^{t-1}\right) \right| &\leqslant \pi(y) \cdot \left| \sum_{k=2}^{n-1} u_k(x) \cdot u_k(y) \cdot \lambda_k^t \right| \\
&\leqslant \frac{1}{n} \cdot \max_{2 \leqslant k \leqslant n-1} \left|\lambda_k^t\right| \cdot \frac{1}{n} \cdot \sum_{k=2}^{n-1} |u_k(x) \cdot u_k(y)| \\
&\leqslant \frac{1}{n} \cdot \max_{2 \leqslant k \leqslant n-1} \left|\lambda_k^t\right| \cdot \sqrt{\sum_{k=2}^{n-1} u_k(x)^2 \cdot \sum_{k=2}^{n-1} u_k(y)^2}
\end{aligned}$$

As in [23, Proof of Theorem 12.3], using the orthonormality of the eigenvectors, we have

$$\sum_{k=2}^{n-1} u_k(x)^2 \leqslant \sum_{k=2}^{n} u_k(x)^2 \leqslant n,$$

and the second statement follows if u and v are in the same partition. The case where u and v are in different partitions follows analogously.

For the third statement, first note that by [23, Exercise 12.3], all eigenvalues of the transition matrix M are non-negative. Since all eigenvalues are non-negative, we conclude from (20) that $p^t_{x,x}$ is non-increasing in t as needed. ∎

The following is a simple corollary from a recent work by Marcus et al. [26] on the existence of Ramanujan graphs.

Lemma A.2 (cf. Marcus et al. [26]). *For any integer $d \geqslant 3$, there are d-regular graphs $H = (V, E)$ with $t_{\text{mix}} = O(\log n / \log(d))$.*

Proof. Marcus et al. [26] show that the existence of a d-regular bipartite Ramanujan graph H such that $\max\{\lambda_2(\widehat{Q}), |\lambda_{n-1}(\widehat{Q})|\} = O(1/\sqrt{d})$, where $\widehat{Q} = \frac{1}{d}A$ is the transition matrix of a non-lazy random walk where A is the adjacency matrix. By the second statement of Lemma A.1, for any pair of states x, y in the same partition

$$p^t_{x,y} \leqslant \frac{1}{n} \cdot \left(1 + (-1)^{t-1}\right) + \max\{\lambda_2, |\lambda_{n-1}|\}.$$

Similarly, x and y are in opposite partitions,

$$p^t_{x,y} \leqslant \frac{1}{n} \cdot \left(1 + (-1)^{t}\right) + \max\{\lambda_2, |\lambda_{n-1}|\}.$$

To translate from the non-lazy random walk $\widehat{Q}$ to a lazy-random walk P, let Z denote the number of non-loops performed by a lazy random walk of length s. We choose s to be the smallest odd integer being greater than $\max\{100, 10t\}$,

$$\mathbb{P}\left[\, Z \geqslant t \,\right] \geqslant 9/10,$$

and by symmetry and the fact that s is odd, $\mathbb{P}\left[\, Z \text{ is even} \,\right] = \frac{1}{2}$. Hence by the Union bound,

$$\mathbb{P}\left[\, Z \text{ is even} \mid Z \geqslant t \,\right] \geqslant \mathbb{P}\left[\, Z \text{ is even} \cap Z \geqslant t \,\right] \geqslant \mathbb{P}\left[\, Z \text{ is even} \,\right] - \mathbb{P}\left[\, Z < t \,\right] \geqslant \frac{4}{10},$$

and similarly, $\mathbb{P}\left[\, Z \text{ is odd} \mid Z \geqslant t \,\right] \geqslant \frac{4}{10}$, and finally,

$$\|p^s_{u,\cdot} - \pi\|_{\text{TV}} \leqslant \mathbb{P}\left[\, Z < t \,\right] \cdot 1 + \mathbb{P}\left[\, Z \geqslant t \,\right] \cdot \left(\sum_{v \in V} \left| \frac{6}{10} \cdot 2/n - 1/n \right| \right) \leqslant \frac{1}{10} + \frac{1}{5} < 1/e,$$

where the first inequality follows from the equations for $p^t_{x,y}$ above. □

The following lemma will be helpful to define a coupling between distributions that are close to the stationary distribution and the exact stationary distribution. (A very similar lemma has been derived in [19, Lemma 2.8])

Lemma A.3. *Let $\varepsilon \in (0, 1]$ be an arbitrary value. Let Z_1 and Z_2 be two probability distributions over $\{1, \ldots, n\}$ so that $\mathbb{P}\left[\, Z_1 = i \,\right] \geqslant \varepsilon \cdot \mathbb{P}\left[\, Z_2 = i \,\right]$ for every $1 \leqslant i \leqslant n$. Then, there is a coupling $(\widetilde{Z}_1, \widetilde{Z}_2)$ of (Z_1, Z_2) and an event $\mathcal{E}$ with $\mathbb{P}\left[\, \mathcal{E} \,\right] \geqslant \varepsilon$ so that*

$$\mathbb{P}\left[\, \widetilde{Z}_1 = i \mid \mathcal{E} \,\right] = \mathbb{P}\left[\, \widetilde{Z}_2 = i \,\right] \qquad \text{for every } 1 \leqslant i \leqslant n.$$

Proof. Let $U \in [0, 1]$ be a uniform random variable. We next define our coupling $(\widetilde{Z}_1, \widetilde{Z}_2)$ of Z_1 and Z_2 that will depend on the outcome of U. First, if $U \in [0, \varepsilon)$, then we set

$$\widetilde{Z}_1 = \widetilde{Z}_2 = i, \qquad \text{if } i \text{ satisfies } \varepsilon \textstyle\sum_{k=1}^{i-1} \mathbb{P}\left[\, Z_2 = k \,\right] \leqslant U < \varepsilon \sum_{k=1}^{i} \mathbb{P}\left[\, Z_2 = k \,\right].$$

For the case where $U \in (\varepsilon, 1)$, it is clear that the definition of U can be extended in a way so that $\widetilde{Z}_1$ has the same distribution as Z_1, and $\widetilde{Z}_2$ has the same distribution as Z_2. Furthermore, notice that if $U \in [0, \varepsilon)$ happens, then $\widetilde{Z}_1$ has the same distribution as Z_2, and $\widetilde{Z}_1 = \widetilde{Z}_2$. Observing that $\mathbb{P}\left[\, U \in [0, \varepsilon) \,\right] = \varepsilon$ completes the proof. □

The following lemma is an immediate consequence of Lemma A.3.

Lemma A.4. *Consider a random walk $(X_t)_{t\geqslant 0}$, starting from an arbitrary but fixed vertex x_0. Then with probability at least $1-1/e$, we can couple $X_{4t_{\mathrm{mix}}}$ with the stationary distribution.*

Proof. Consider the random walk $(X_t)_{t\geqslant 0}$ after step $s := t_{\mathrm{sep}} \leqslant 4t_{\mathrm{mix}}$. By definition of t_{sep}, $p^s_{u,v} \geqslant (1-1/e)\pi(v)$. Applying Lemma A.3, where Z_1 is the distribution given by $p^t_{u,v}$ and Z_2 is the stationary distribution shows that with probability at least $1-1/e$, X_s has the same distribution as π. If this is the case, then the same holds for $X_{4t_{\mathrm{mix}}}$ as well. □

The lemma above shows that for t_{meet} and t_{coal} it suffices to consider the stationary case:

Lemma A.5. *For any graph G,*

$$\max\{(1/e)t_{\mathrm{mix}}, t^{\pi}_{\mathrm{meet}}\} \leqslant t_{\mathrm{meet}} \leqslant \frac{2}{(1-1/e)^2}\cdot(4t_{\mathrm{mix}} + 2t^{\pi}_{\mathrm{meet}})\,,$$

and similarly, $t_{\mathrm{coal}} \leqslant 4\cdot(4t_{\mathrm{mix}} + 2t^{\pi}_{\mathrm{coal}})$.

Proof. We begin by proving the lower bound on t_{meet}. First, consider two independent random walks $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$ that are run for $t = e\cdot t_{\mathrm{meet}}$ time-steps. Then, we have

$$\bar{d}(t) = \max_{u,v} \|p^t_{u,\cdot} - p^t_{v,\cdot}\|_{\mathrm{TV}} \leqslant \mathbb{P}\left[\cup_{s=0}^t X_s = Y_s\right] \leqslant \frac{1}{e},$$

where the first inequality is due to the coupling method [23, Theorem 5.3] and the second inequality follows by Markov's inequality. The above inequality implies $t_{\mathrm{mix}} \leqslant e\cdot t_{\mathrm{meet}}$. Furthermore, $t^{\pi}_{\mathrm{meet}} \leqslant t_{\mathrm{meet}}$ holds by definition, and the lower bound follows.

For the upper bound, we divide the two random walks into consecutive epochs of length $\ell := 4t_{\mathrm{mix}} + 2t^{\pi}_{\mathrm{meet}}$. For the statement it suffices to prove that in each such epoch, regardless of the start vertices of the two random walks, a meeting occurs with probability at least $(1-1/e)^2\cdot 1/2$.

Consider the first random walk $(X_t)_{t\geqslant 0}$ starting from an arbitrary vertex after $s := 4t_{\mathrm{mix}}$ steps. By Lemma A.4, we obtain that with probability at least $1-1/e$, the distribution of X_s is equal to that of a stationary random walk. Similarly, we obtain that with probability at least $1-1/e$, the distribution of Y_s is equal to that of a stationary distribution. Hence with probability $(1-1/e)^2$, X_s and Y_s are drawn independently from the stationary distribution. In this case, it follows by Markov's inequality that the two random walks meet before step $s + 2t^{\pi}_{\mathrm{meet}}$ with probability at least $1/2$. Overall, we have shown that with probability at least $(1-1/e)^2\cdot 1/2$, a meeting occurs in a single epoch. Since this lower bound holds for every epoch, independent of the outcomes in previous epochs, the upper bound on the expected time t_{meet} follows.

The upper bound on t_{coal} in terms of t^{π}_{coal} is shown in exactly the same way. □

Observation A.6. *Consider two non-interacting random walks $(X_t)_{t\geqslant 0}$ and $(Y_t)_{t\geqslant 0}$ starting on nodes drawn from the stationary distribution. Fix an arbitrary $t \in \mathbb{N}$. Define the collision-counting random variables $Z_1 = \sum_{i=0}^{\lceil t/2\rceil} \mathbf{1}_{X_t=Y_t}$, $Z_2 = \sum_{i=\lceil t/2\rceil+1}^{t} \mathbf{1}_{X_t=Y_t}$, and $Z = Z_1 + Z_2$. Then $\mathbb{P}[Z_1 \geqslant 1 \mid Z \geqslant 1] \geqslant \frac{1}{2}$.*

Proof. Since both nodes start from the stationary distribution, $\mathbb{P}[Z_1 \geqslant 1] \geqslant \mathbb{P}[Z_2 \geqslant 1]$. By the Union bound, $\mathbb{P}[Z \geqslant 1] \leqslant 2\cdot\mathbb{P}[Z_1 \geqslant 1]$. By law of total probability, $\mathbb{P}[Z_1] = \mathbb{P}[Z_1 \geqslant 1 \mid Z \geqslant 1]\cdot \mathbb{P}[Z \geqslant 1]$. Putting everything together yields $\mathbb{P}[Z_1 \geqslant 1 \mid Z \geqslant 1] \geqslant \frac{1}{2}$. □

Lemma A.7. *Let $(X_t)_{t\geqslant 0}$ be a stochastic process satisfying (i) $\mathbb{E}[\,X_t|\mathcal{F}_{t-1}\,] \leqslant \beta \cdot X_{t-1}$, for some $\beta < 1$, and (ii) $X_t \geqslant 0$ for all $t \geqslant 0$. Let $\tau(g) = \min\{t \geqslant 0 | X_t \leqslant g\}$ for $g \in (0, |X_0|)$, then*

$$\mathbb{E}[\,\tau(g)\,] \leqslant 2 \cdot \lceil \log_\beta(g/(2X_0)) \rceil.$$

Proof. By the iterative law of expectation, we have

$$\mathbb{E}[\,X_t\,] \leqslant \beta^t \cdot X_0.$$

Furthermore, by Markov's inequality, for any $\lambda \geqslant 1$

$$\mathbb{P}\left[\,\tau(g) > \lambda \cdot \lceil \log_\beta(g/(2X_0)) \rceil\,\right] \leqslant \mathbb{P}\left[\,X_{\lambda \cdot \lceil \log_\beta(g/(2X_0)) \rceil} > g\,\right] \leqslant \frac{\mathbb{E}\left[\,X_{\lambda \cdot \lceil \log_\beta(g/(2X_0)) \rceil}\,\right]}{g} \leqslant 2^{-\lambda}.$$

Therefore,

$$\begin{aligned}
\mathbb{E}[\,\tau(g)\,] &= \sum_{i=1}^{\infty} \mathbb{P}[\,\tau(g) \geqslant i\,] \\
&\leqslant \lceil \log_\beta(g/(2X_0)) \rceil + \sum_{\lambda=1}^{\infty} \lceil \log_\beta(g/(2X_0)) \rceil \cdot \mathbb{P}\left[\,\tau(g) > \lambda \cdot \lceil \log_\beta(g/(2X_0)) \rceil\,\right] \\
&\leqslant 2 \cdot \lceil \log_\beta(g/(2X_0)) \rceil.
\end{aligned}$$

□